NAVAL POSTGRADUATE SCHOOL
Monterey, California

THESIS

REPRESENTING TACTICAL LAND NAVIGATION EXPERTISE

by

Jason L. Stine

September 2000

Thesis Advisor: Rudolph Darken
Thesis Co-Advisor: Barry Peterson

Approved for public release; distribution is unlimited.

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE September 2000	3. REPORT TYPE AND DATES COVERED Master's Thesis

4. TITLE AND SUBTITLE Representing Tactical Land Navigation Expertise	5. FUNDING NUMBERS
6. AUTHOR(S) Jason L. Stine	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research 800 N. Quincy St. Tower 1 Arlington, Virginia 22217-5660	10. SPONSORING / MONITORING AGENCY REPORT NUMBER
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.	
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.	12b. DISTRIBUTION CODE

13. ABSTRACT *(maximum 200 words)*

Tactical land navigation is a very important, but extremely difficult task performed daily by small unit leaders. In an effort to find ways to develop expertise more efficiently, a detailed description of expert performance is presented and contrasted with novice and intermediate performance. This definition fits the Recognition Primed Decision model of human cognitive behavior. Then, through use of the Critical Decision Method of knowledge elicitation, interviews with experts at the U. S. Army Special Forces Qualification Course formed the basis of a detailed cognitive model of expert tactical land navigation. Four important characteristics of experts emerge: (1) they rely on high-fidelity mental maps; (2) they blend multiple cues; (3) they adjust and recalibrate tools dynamically; and (4) they visualize spatial information. Finally, a multi-agent system computationally represents the route planning portion of the performance model.

14. SUBJECT TERMS Agent Based Modeling, Land Navigation, Multi-Agent System, Human Performance Modeling			15. NUMBER OF PAGES 102
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFI-CATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited.

REPRESENTING TACTICAL LAND NAVIGATION EXPERTISE

Jason L. Stine
Major, United States Army
B.S., United States Military Academy, 1989

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN MODELING, VIRTUAL ENVIRONMENTS, AND SIMULATION

from the

NAVAL POSTGRADUATE SCHOOL
September 2000

Author: ______________________
Jason L. Stine

Approved by: ______________________
Rudolph Darken, Thesis Advisor

____
Barry Peterson, Co-Advisor

____
Rudolph Darken, Academic Associate
Modeling, Virtual Environments, and Simulation
Academic Group

____
Michael Zyda, Chair
Modeling, Virtual Environments, and Simulation
Academic Group

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Tactical land navigation is a very important, but extremely difficult task performed daily by small unit leaders. In an effort to find ways to develop expertise more efficiently, a detailed description of expert performance is presented and contrasted with novice and intermediate performance. This definition fits the Recognition Primed Decision model of human cognitive behavior. Then, through use of the Critical Decision Method of knowledge elicitation, interviews with experts at the U. S. Army Special Forces Qualification Course formed the basis of a detailed cognitive model of expert tactical land navigation. Four important characteristics of experts emerge: (1) they rely on high-fidelity mental maps; (2) they blend multiple cues; (3) they adjust and recalibrate tools dynamically; and (4) they visualize spatial information. Finally, a multi-agent system computationally represents the route planning portion of the performance model.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I. INTRODUCTION 1
A. PROBLEM STATEMENT 1
B. MOTIVATION 1
1. Tactical Navigation vs. Sport Orienteering 1
2. Cognitive, then Computational Model 2
3. Army and DoD Relevance 3
C. THESIS ORGANIZATION 3
II. BACKGROUND AND PREVIOUS WORK 5
A. NATURALISTIC DECISION MAKING 5
1. NDM Characteristics 5
2. Dreyfus & Dreyfus - 5 Levels of Proficiency 6
3. Recognition Primed Decision Model 9
4. RPD and Expertise 13
B. CRITICAL DECISION METHOD 14
C. ADAPTIVE, AUTONOMOUS AGENTS 14
1. Multi-Agent Systems 14
2. Irreducible Semi-Autonomous Adaptive Combat 17
3. Genetic Algorithms 20
III. COGNITIVE TASK ANALYSIS: TOOLS AND MODELS 23
A. MODELING EXPERT NAVIGATORS USING NDM 23
B. DEFINING EXPERTISE 25
1. Exploratory Focus Group 25
2. Differences Between Experts, Competents, and Novices 25
a. Planning Phase 25
b. Patrol Base to ORP Phase 28
c. ORP to Objective Phase 30
d. Discussion 31
3. Identification of Experts 32
C. KNOWLEDGE ELICITATION USING CDM 33
1. Planned vs. Actual Protocols 34
2. Setting and Conduct of Interviews 35
IV. REPRESENTATION OF EXPERTISE 37
A. GENERAL RESULTS 37
1. General 37
2. Mental Map Fidelity 37
3. Blending of Cues 38
4. Dynamic Adjustment 39
5. Spatial Visualization 40
B. CUE INVENTORY 41

1. General ... 41
2. Navigation Tools ... 41
3. Environmental Conditions ... 42
4. Mission Conditions ... 43
5. Terrain Features ... 43
C. SITUATION ASSESSMENT RECORD ... 45
1. General ... 45
2. Situation Assessment Record, Variation 1 ... 48
3. Situation Assessment Record, Variation 2 ... 49
4. Situation Assessment Record, Variation 3 ... 52
D. KEY DECISION REQUIREMENTS ... 53

V. ROUTE PLANNING EXECUTABLE MODEL ... 55

A. AGENT-BASED ROUTE PLANNER ... 55
1. Focus on Route Planning ... 55
2. Selection of an Agent-based Representation ... 56
B. TERRAIN REPRESENTATION ... 56
1. Organization of OpenFlight Files ... 56
2. User Interface ... 57
3. Enemy Locations ... 58
C. AGENT DESIGN AND STRUCTURE ... 58
1. Decision Vector ... 58
2. Perception Vector ... 59
3. Establishing Checkpoints ... 60
4. Movement Using Lowest Penalty Function ... 61
D. SCORING AND GENETIC ALGORITHM ... 63
1. Multiple Agents and Scoring ... 63
2. Genetic Algorithm ... 66
3. Achieving Expertise ... 67

VI. CONCLUSIONS AND RECOMMENDATIONS ... 69

A. APPLICATION OF THE RPD MODEL ... 69
1. CDM Implementation ... 69
2. Identifying Experts ... 69
B. EXECUTABLE REPRESENTATIONS ... 70
1. Suitability of Agent-based Models ... 70
2. Terrain Representation Techniques ... 71
3. Portability ... 71
C. TRAINING STRATEGIES ... 71
1. Training vs. Experience ... 71
2. Use of Computer Based Trainers ... 73
3. Other Possible Trainers ... 73
D. FUTURE WORK ... 74
1. Cognitive Task Analysis ... 74
2. Improved Agent-based Route Planner ... 74

APPENDIX A. SMALL UNIT MISSIONS **77**
1. GENERAL INFORMATION 77
2. PLANNING 77
3. MOVEMENT 78
4. ACTIONS ON THE OBJECTIVE 79

LIST OF REFERENCES **81**

INITIAL DISTRIBUTION LIST **83**

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

Figure 1. Simple Match from (Klein, 1998).. 10
Figure 2. Diagnose the Situation from (Klein, 1998).. 11
Figure 3. Evaluate a Course of Action from (Klein, 1998).. 12
Figure 4. Two-dimensional ISAAC Battlefield from (Ilachinski, 1997).. 18
Figure 5. ISAAC Penalty Function from (Ilachinski, 1997).. 19
Figure 6. Cue Inventory grouped by Category.. 41
Figure 7. Specific By-products of Land Navigation Situation Assessment from (Klein, 1998).. 48
Figure 8. Situation Assessment Record for RPD Variation 1.. 49
Figure 9. Situation Assessment Record for RPD Variation 2.. 51
Figure 10. Situation Assessment Record for RPD Variation 3.. 53
Figure 11. Key Decision Requirements from (Hoffman, et al., 1995).. 54
Figure 12. Sample Penalty Function Calculation.. 62
Figure 13. Sample Scoring Calculation.. 65

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS

AI	Artificial Intelligence
CDM	Critical Decision Method
CGF	Computer Generated Forces
DTED	Digital Terrain Elevation Data
ISAAC	Irreducible Semi-Autonomous Adaptive Combat
NDM	Naturalistic Decision Making
NIMA	National Imagery and Mapping Agency
ORP	Objective Rally Point
PL	Patrol Leader
RPD	Recognition Primed Decision
RTB	Ranger Training Battalion
VRML	Virtual Reality Modeling Language

I. INTRODUCTION

A. PROBLEM STATEMENT

The purpose of this thesis is to present a detailed description of an expert dismounted land navigator in a tactical military setting, then to show how a portion of this expertise may be represented in an executable, personal computer-based model. From this model, future work may develop a more complete model of expert navigation.

B. MOTIVATION

1. Tactical Navigation vs. Sport Orienteering

One of the most embarrassing, and most common, experiences for a new infantry leader is to get his unit lost. Tactical land navigation is a complicated and difficult task, and it is inherently different from civilian orienteering. The tactical navigator incorporates elements such as small unit tactics, group leadership, and military mission planning that do not exist for the civilian orienteer. Novices become experts through repetition, which is an often long and painful experience.

Currently, the Army conducts novice land navigation training on civilian-style orienteering courses. This allows efficient use of training time, since many soldiers can train basic navigation tasks at the same time. Tactical navigation, however, is a related, but distinctly different task. While sport and tactical navigation do rely on a set of common skills, an experienced tactical navigator uses a specific tactical strategy. Some elements of that tactical strategy may transfer directly from a sport navigation task, but other elements must be adapted to create a different sport strategy. Because the two tasks

require different types of strategies, proficiency in either sport or tactical navigation does not necessarily transfer into the other.

2. Cognitive, then Computational Model

A virtual environment based training device that improves a novice infantry leader's tactical navigation skills would be invaluable to the Army. An ideal trainer would both increase confidence and shorten the development time from novice to expert. However, to properly train the attributes of expertise, they must first be defined and quantified. Navigation is by nature a highly aggregated cognitive task, making it very difficult to break down and measure. A detailed cognitive model will define the important tasks that the expert performs well, the cues he needs to make decisions, and the expert's reactions to unexpected situations. Without such information, a realistic computational model of expert performance is impossible.

An ideal virtual environment training system supplements, rather than replaces, traditional, physical navigation practice with targeted virtual environment practice sessions. One version of such virtual training provides feedback to the student while practicing in the virtual environment via a computer generated virtual expert navigator, which drives the feedback module.

The executable model must be based on the cognitive model. To perform realistically, the model should process the same cues and patterns and react to these cues in the same ways as human experts. The scope of expert tactical military navigation is very large, and representing it in its entirety is beyond the scope of this thesis. As discussed later, these scope and implementation factors influenced the immediate

computational model, which focuses on the route planning component of expert performance.

3. Army and DoD Relevance

Two main application areas exist: training tools and realistic computer generated forces (CGF). An executable model could be used as an expert "tutor" to teach less experienced navigators to plan routes. This would free the human trainers to concentrate on training the real task in a tactical environment. Also, the model could aid in developing more realistic models of small-scale military operations. For example, if an analyst wished to evaluate the possible consequences of using an elite squad or a conventional squad for a given mission, he might use this model. The elite squad would use the expert route planner, reflecting the higher level of navigation expertise in the unit, and the conventional squad would use a less proficient version that plans intermediate routes. Both expert and intermediate models can be produced using the system presented in Chapter VI. In another application, the expert route planner could be used to control the movement of elite enemy forces, or enemy forces that are operating on familiar terrain. In short, any scenario that calls for expert navigation behavior could use this model.

C. THESIS ORGANIZATION

This thesis is organized in the following manner: Chapter II explores the background of Naturalistic Decision Making theory, the Recognition Primed Decision model of expert behavior, the Critical Decision Method of knowledge elicitation, and the use of adaptive, autonomous software agents to solve ill-defined problems. Chapter III outlines the selection of a cognitive modeling architecture and collection of data. Chapter

IV describes the cognitive model in detail. Chapter V discusses the agent based executable model of expert route planning. Chapter VI provides the conclusions and the recommended areas for future research.

There is one appendix that provides an explanation of small unit military operations.

II. BACKGROUND AND PREVIOUS WORK

A. NATURALISTIC DECISION MAKING

1. NDM Characteristics

Traditionally, the most widely accepted explanation of human decision making is the Rational Choice Theory. According to this theory, people explicitly deliberate between possible alternatives, then select the best course of action (Zsambok, 1997). Several field studies in complex work environments indicate that experienced decision makers behave in a very different manner. Experienced people use experience to form mental simulations that suggest solutions to difficult problems. This allows such complex decisions to be made quickly, without decomposing problems into smaller elements, which can be analyzed. Gary Klein and Caroline Zsambok, in an attempt to explain this behavior, began to study what is now called Naturalistic Decision Making (NDM). They focus on how people use their knowledge and experience to assess complex, uncertain conditions in real, uncontrolled situations, and then take action.

Klein and Zsambok define four key elements of NDM. First, the task and setting involve ill-structured problems, dynamic environments, competing goals, high stakes, time pressure, and organizational goals and norms. Thus, traditional, carefully controlled scientific experiments are difficult or impossible to conduct because the conditions are numerous, intertwined, and cannot be duplicated. Second, the subjects are experienced participants with extensive knowledge of the task. Tracing the use of this experience is one of the key differences between the NDM approach and classical models of decision making. Third, the moment the subject makes his decision is not the important point: rather, NDM research focuses on the inputs to the decision: situation awareness and plan

generation. Fourth, the purpose of research is to describe expert strategies rather than prescribe the strategies everyone across all ability levels ought to use (Zsambok and Klein, 1997).

NDM is useful primarily because it embraces the complex environments that are frequently encountered in the real world but never duplicated in any laboratory. It offers a way to explain how real people solve real problems, and how experienced decision makers differ from novices. Moreover, it encompasses group dynamics and organizational culture, important elements of the real world that most classical models try to minimize. Thus, it provides a good framework for study of decision making in military settings (Zsambok and Klein, 1997), as well as that in police and fire departments, Fortune 500 companies, and political parties. However, useful research depends on the ability to identify and access expert subjects; this is certainly not a trivial task.

2. Dreyfus & Dreyfus - 5 Levels of Proficiency

Hubert and Stuart Dreyfus described the differences between individuals with different levels of experience and competence with a five stage model of skill acquisition (Dreyfus, 1997). They describe their theory in terms of cars and drivers and chess players to illustrate the ideas in terms of both motor skills and intellectual skills.

The first stage, novice, deals with beginners and their instructors. The instructor decomposes the task into simple elements separated from any real-world situation, each of which the beginner can recognize even without experience. Then, the instructor supplies the beginner with a set of rules for determining actions based on the state of these simple, context-free elements, in the same way that a computer executes a program.

For example, student drivers learn to shift to second gear when the speedometer needle reaches ten miles per hour. Novice chess players learn numerical values for each piece and the general rule to exchange pieces with the opponent if the total value of pieces captured is greater than that of pieces lost.

As novices gain experience, they progress to the second stage, advanced beginner. Advanced beginners recognize important aspects of real world situations that help to make better decisions. Thus, they begin to use situational aspects as well as the situation independent rules they learned as novices. Advanced beginner drivers learn to shift based on the sound of the engine without looking at the speedometer needle. Advanced beginner chess players learn to recognize unfavorable positions and how to avoid them.

A hierarchical perspective of the important factors of the situation characterizes stage three, competence. The advanced beginner begins to realize that the number of potentially important cues in any situation is overwhelming, so he develops a plan that then determines which elements are important and which can be ignored. Thus, competent performers look to develop rules to decide on which plan or perspective is appropriate. Unfortunately, there are a very large number of possible situations, which differ in subtle ways, and no performer could possibly memorize a list of rules for each possibility in the same fashion as novices memorize their rules. Therefore, competent performers must decide on the best available plan without knowing for sure whether or not that plan will work. Performance of the task is, for the competent, a nerve-wracking and often frightening experience, since he has enough experience to understand the risks but not enough to guarantee success. In other words, he has outgrown the simple novice rule sets, but has not yet effectively replaced them with something better. For example,

competent drivers take into account speed, surface condition, and space available to decide if the car is moving too quickly. Then, they must decide whether to let up on the accelerator or step on the brake, and are happy to get through curves without mishap. A competent chess player may decide, after studying the board, that attacking the opposing king is the best option. He will ignore indications of weaknesses in his own position created by the attack, and so may be vulnerable to counterattack.

Most performers do not achieve stage four - proficiency. Those who become proficient manage to integrate experience into the theory of skill, and replace the system of rules and responses with a set of situational discriminations and associated responses. In short, behavior becomes less deliberate and more intuitive. Action is easier and less stressful because the performer simply sees what the appropriate goal is in each situation, instead of having to decide by calculating the value of several alternatives. Proficient drivers, for example, know intuitively when the car is going too fast. They still must decide which action is appropriate, but they save valuable time by avoiding the decision of whether or not speed is excessive. Proficient chess players, classified as masters, recognize a large number of types of positions and can recognize these without conscious effort. They know immediately that, in a given situation, attack is the proper course of action, but must deliberate about how to do it best.

What distinguishes an expert (stage five) from a proficient performer is a more refined discrimination ability that allows him to intuitively recognize not only what should be done but also how to do it. Enough experience in a variety of situations allows the performer to group situations into classes which share the same decision, action, or tactic. This allows the immediate intuitive responses characteristic of, and limited to,

expertise. Expert drivers not only know that the car is going too fast, but also respond with their feet appropriately on the accelerator or brake. Expert chess players, classified as international grandmasters, recognize up to 50,000 types of positions and can play at a rate of 5 to 10 seconds a move without any degradation in skill (Dreyfus. 1997). Experts, however, may sometimes revert to competent performance if confronted with totally novel situations.

Clearly, only a small percentage of performers ever reach the expert level, at which performance is almost exclusively intuitive. Rational choice theory requires explicit deliberation of several possible options to explain performance. It does not account for the intuitive aspects of proficient and expert performance. Rational choice theory, while a good way to describe novice through competent performance, is then clearly unsuitable for a detailed cognitive analysis of proficient and expert performance.

3. Recognition Primed Decision Model

NDM theorist Gary Klein developed the Recognition Primed Decision (RPD) model, which attempts to explain how experienced decision makers use their expertise to identify and carry out a course of action without having to analyze several options for the purpose of comparison (Klein, 1998). The model contains three functions, one of which the expert will apply when faced with a decision.

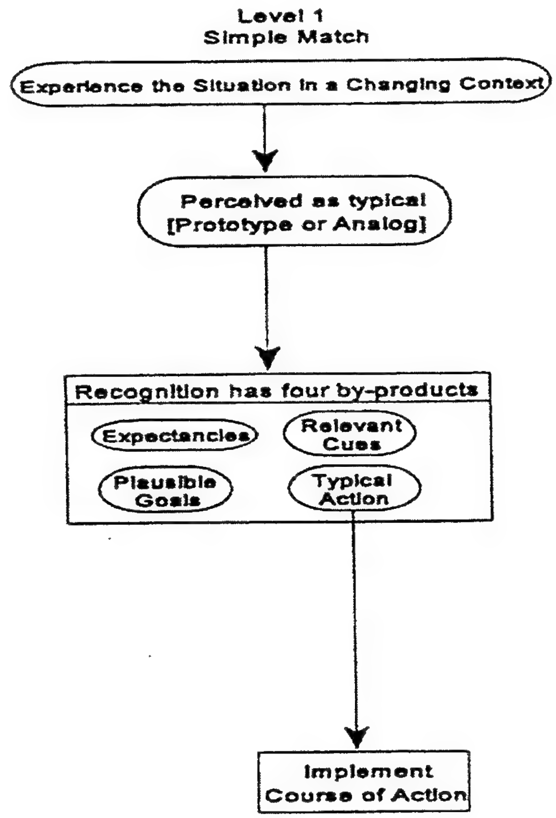

Figure 1. Simple Match from (Klein, 1998)

The "Simple Match" function represents a straightforward case in which the decision maker identifies a situation and reacts accordingly, completely without deliberation. The goals are clear, critical cues are recognizable and within expected parameters, and a typical course of action is readily apparent. In this case, the expert may not even recognize that he made a decision.

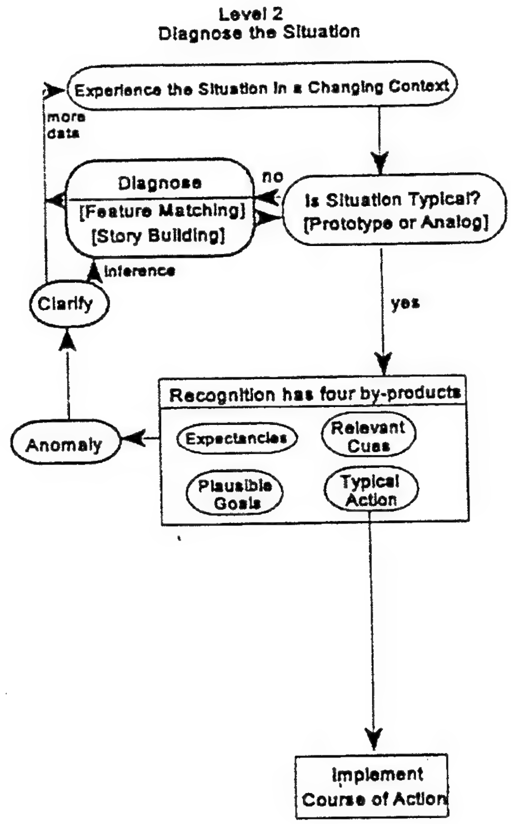

Figure 2. Diagnose the Situation from (Klein, 1998)

A more complicated or unfamiliar situation forces the expert to use the "Diagnose the Situation" function. This function allows the expert to link observed events to causal factors, thereby explaining the events and allowing the expert to classify them according to his experience, and intuitively generate the appropriate response. The expert may either try to match the features of observed events to those of situations with which he is familiar (feature matching) or generate a new explanation of causes for the observed effects (story building) which would also fit with his experience.

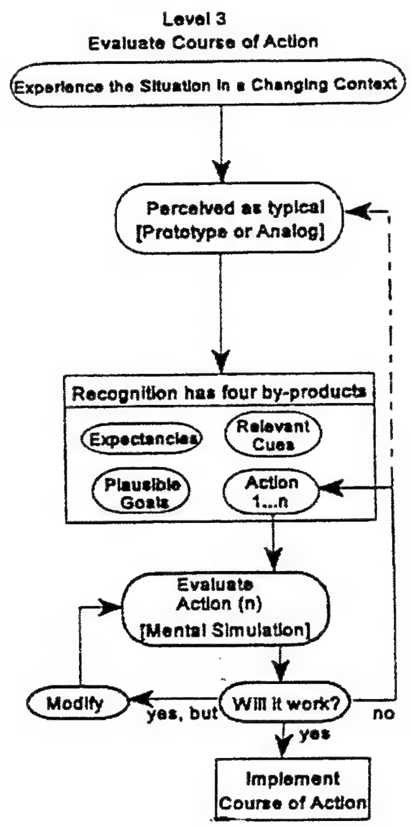

Figure 3. Evaluate a Course of Action from (Klein, 1998)

The most complex case results in the expert's use of the “Evaluate a Course of Action” function. In this case, the situation is new and unfamiliar, so the expert develops a mental simulation of his intuitive response. He determines whether or not the course of action will run into difficulties, whether or not these difficulties can be remedied, or whether a new course of action will be required.

The RPD model has been used to explain the performance of experts in a wide variety of activities, including urban and rural firefighting, flight control in commercial airlines, chess tournament play, and intensive care unit nursing. In a study of chess players, Klein, Wolf, Militello, and Zsambok (1995) provided empirical data to support

three key assertions of the RPD model: first, that experienced performers will generate a plausible option as the first one they consider; second, that time pressure does not cripple the performance of experts as it does less experienced performers, because experts use pattern matching to make decisions quickly; third, that experts adopt courses of action without comparing and contrasting possible alternatives. In each case, the researchers found statistically significant differences between expert decision makers and competent ones.

4. RPD and Expertise

Although expert performers can operate totally intuitively, they may not always do so. Even experts sometimes make mistakes in complicated environments. Moreover, totally unfamiliar environments may rob the expert of his ability to act intuitively. In these cases, the RPD model describes how experts respond.

Variation 1 of the RPD model describes an expert exclusively using intuition for both situational awareness and decision-making. Everything proceeds according to plan, and the expert's experience covers every situation. In familiar environments, this is the normal mode of operation for experts.

Variation 2 marks a departure from totally intuitive action. Normally, an expert uses this technique when faced with a new, unfamiliar environment. This is almost the inverse of proficient behavior, since it couples deductive situation identification with intuitive action.

Variation 3 explains the expert's actions when he makes a mistake or is confused by an environment radically different from that he expected. It may seem reminiscent of competent behavior, with deductive situation identification followed by deductive action.

However, the mental simulation and course of action adjustments performed by experts are clearly beyond the capabilities of competent performers.

B. CRITICAL DECISION METHOD

To describe expert behavior using the RPD model, the researcher must conduct a cognitive task analysis to discern what experts actually do when faced with real world problems. NDM research suggests using a knowledge elicitation technique called the Critical Decision Method (CDM)

At the heart of this method is the critical decision itself (Klein, Calderwood, and MacGregor, 1989). The interview is structured to first identify an appropriately critical decision and then probe deeply into the cognitive operations that resulted in the decision. As described by Hoffman, Shadbolt, Burton and Klein (1995), the procedure is structured as follows. During interview preparation, the elicitor learns about the domain and gains access to the participants. Once the specific interview begins, the first step is to select the incident; the elicitor works with the participant to identify a situation in which the expert's skills were challenged and it stands out in the decision-maker's mind as being critical. Once identified, the elicitor guides the participant through progressively deeper and more detailed retellings of the incident. Typically, the interview is audio or video recorded.

C. ADAPTIVE, AUTONOMOUS AGENTS

1. Multi-Agent Systems

Given that Dreyfus (1997) defined novice behavior as rule-based and similar to a computer following a program, it stands to reason that traditional rule-based artificial intelligence (AI) approaches are suited to novice performance. Since experts do not

operate in a deductive, rule-based fashion, however, traditional methods are unsuited to realistic models of expert behavior. A different software architecture, that of a multi-agent system, is better suited to computational representations of the RPD model.

Jacques Ferber (1999) defines a software agent as a construct that has several important capabilities. First, it is capable of acting, not just reasoning as traditional AI constructs do. They carry out actions that will change the environment in which they operate, thus affecting future decisions. Second, they are autonomous, meaning that they are not controlled by the user but act in accordance with a set of tendencies. These tendencies may represent individual goals or satisfaction/survival functions, which the agent attempts to optimize. Thus, the agent can accept or reject requests from other agents, and has some freedom of movement, which allows it to reject certain goals or rules in certain situations. Third, agents can perceive the environment in which they operate, but only to a limited extent. They lack the global knowledge of the situation common to most AI constructs, and so receive information in a manner more like that of humans. Fourth, agents may have the capacity to reproduce themselves, most often through the use of a genetic algorithm, which reproduces the more successful agents while discouraging reproduction of less successful ones. Given this definition of agents, a multi-agent system is composed of six distinct elements.

First, a multi-agent system is situated in an environment, a physical space that generally has a volume. In most systems, the agent's perception of this physical space is the key factor in its actions. From the agent's perspective, the environment is everything the agent itself is not, so it is impossible to define a situated agent without first defining the environment surrounding it. Second, the system contains a set of objects, each of

which is situated in the environment. These objects are passive in that agents can perceive, create, modify, or destroy them. Third, a multi-agent system must obviously contain a collection of agents that satisfy the above definition. The agents make a subset of the objects and represent the active elements of the system. Fourth, an assembly of relations links objects and agents to each other. These relations define the procedures for communication between entities. Fifth, a set of operations defines the ways in which the collection of agents may perceive, produce, consume, transform, and manipulate members of the set of objects. Finally, rules which Ferber calls the "laws of the universe" (1999) represent the application of the operations and the effects of the operations on the environment. For example, Newtonian physics are the laws of the universe if the environment is a pool table.

The more complex the system, the less likely the programmer trying to represent it will create perfectly tuned agents on the first try. The agent's behavior will initially not resemble human expertise, even with carefully derived satisfaction functions and limited environments. Agents must therefore be able to adapt to be useful or interesting. Many systems allow each agent to modify itself during execution in response to its perceived success or failure in accomplishing its goals. However, the genetic algorithm provides a more controlled framework for adaptation. The genetic algorithm considers each agent's set of tendencies to be its DNA, and agents reproduce by random combinations of two "parent" agents' tendencies (Holland, 1995).

For the genetic algorithm to work properly, the multi-agent system must incorporate a way to assign credit to successful performance and deduct credit from failure. The agents that accumulate the most credit reproduce, while those that fail to

gain enough credit do not reproduce and are removed from the system. As the system runs, more successful agents begin to dominate and more successful behavior is rewarded, while agents with unsuccessful sets of tendencies disappear. Some systems also allow "mutations", or random assignment of tendencies to newborn agents. These tendencies do not appear in either parent agent and fill the same role as mutation in the theory of evolution, introducing new combinations to the species.

2. Irreducible Semi-Autonomous Adaptive Combat

Andrew Ilachinski developed the idea of land combat as a complex adaptive system that could be represented as a multi-agent system. He created Irreducible Semi-Autonomous Adaptive Combat (ISAAC) as a way to explore the group dynamics and adaptive nature of combat (Ilachinski, 1997).

ISAAC contains a situated environment, which in its simplest form is a two-dimensional grid of possible agent locations. The simplest ISAAC contains no objects other than the agents themselves, but more complex versions contain obstacles around which the agents must maneuver. ISAAC agents are roughly analogous to individual combatants. Each time step in the simulation, each agent may move, fire at enemy agents, or stay in its current location. Movement is controlled by a penalty function computed for each possible location to which the agent could move. The location producing the lowest value of the penalty function is the one to which the agent moves.

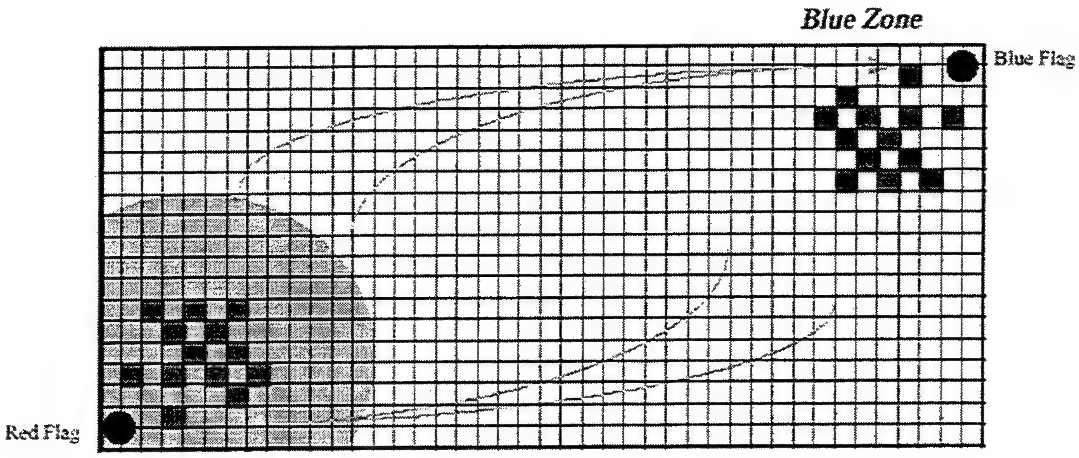

Figure 4. Two-dimensional ISAAC Battlefield from (Ilachinski, 1997)

An ISAAC agent's penalty function consists of six terms, and can be expressed as the scalar product of the agent's tendencies with its perceived environment. The tendencies are to move to the enemy's "flag", to move to the friendly flag, to move toward live enemy agents, to move toward live friendly agents, to move toward wounded enemy agents, and to move toward wounded friendly agents (Ilachinski, 1997). Each agent has a sensor range, within which it knows how many live and wounded friendly and enemy agents are present. Each agent also knows the exact location of its own and the enemy's flags.

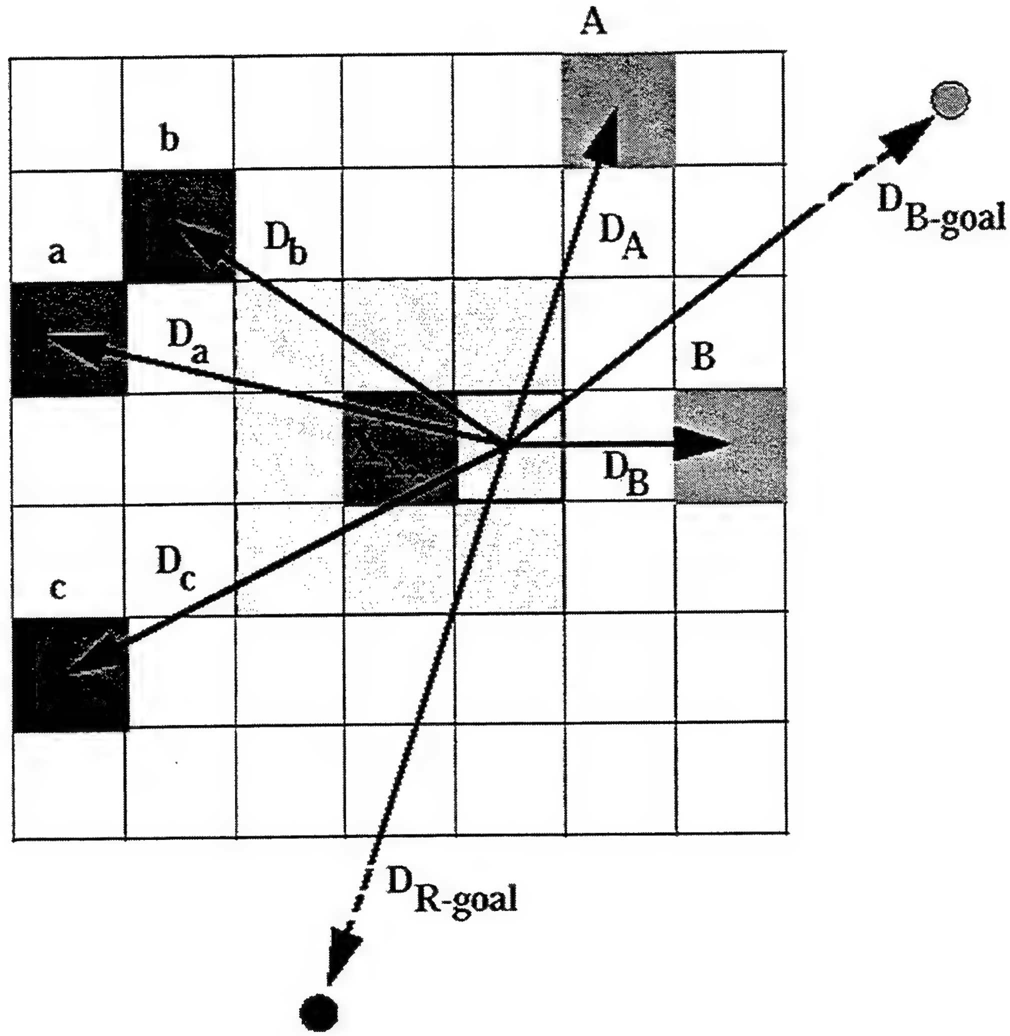

Figure 5. ISAAC Penalty Function from (Ilachinski, 1997)

In Figure 5, a sample penalty calculation is shown. The shaded area indicates the possible spaces to which the agent may move. The penalty function uses the distances between the proposed location (the square from which all the arrows come) and each agent and flag. The calculation for this square represents the value of moving from the current location to the proposed location. The one of the nine possibilities (including the current location) with the lowest value is the square to which the agent moves.

Each agent also has a firing range within which it may engage enemy agents, and wound or kill them according to a random draw compared to a fixed probability of hit. More advanced versions also include commanders who give orders to subordinates and a genetic algorithm to refine each agent's tendencies (Ilachinski, 1997).

3. Genetic Algorithms

A genetic algorithm is a way to find a near-optimal combination of attributes without specifically testing each option (Holland, 1995). Since the possible combinations of agent characteristics in a reasonably complex multi-agent system may be too numerous for reasonable computation, some other way must be found. Even in a relatively small system like ISAAC, there are over 10^{30} possible combinations - an enormous number. Even the fastest computers currently in existence could not test this many cases in a reasonable amount of time. Totally random assignment of characteristics might work, but do not take advantage of earlier trials. Instead, a system of adaptation based on the theory of evolution allows performance improvements in a more reasonable time frame.

Genetic algorithms borrow much terminology from molecular biology and evolutionary theory. They describe agents in terms of parents and children, testing new cases as reproduction, and attributes as genes and alleles. An agent's key decision-making attributes make up its genes, and the different possibilities for each attribute are called alleles. Two agents reproduce by merging these attributes to form the attributes of a totally new agent, the child. Certain attributes may not come from either parent, but may be selected in some other way; these are termed mutations, like their biological equivalents.

Natural selection works by allowing only the "best" or most successful organisms to survive long enough to reproduce. In the natural world, success (from a genetic standpoint) is defined as reproduction, which follows from survival and competition for mates. In the artificial world of the multi-agent system, therefore, there must be a system of credit assignment, or scoring, which allows determination of success and failure. The definition of success is left to the programmer, and will determine the course of evolution in the multi-agent system. The behaviors awarded with credit by the scoring system will endure; those resulting in low scores will disappear.

Evolution takes thousands to millions of years to achieve large-scale changes in the natural world. Computers make the process faster, but still require several generations of agents to refine the desired attributes. Agents must first perform the evaluated action and be scored on their performance. Then, the high-scoring agents are allowed to reproduce, and the low-scoring agents are removed from the system to make room for the offspring. Some systems allow the very best agents to reproduce the most, and moderately successful agents to reproduce to a lesser degree, while some allow only the top performers to reproduce. The new and old agents then perform the task again, and the scoring system again decides which agents will reproduce. Mutations occur, generally at a rate of one every one hundred to one hundred thousand offspring. This quickly can become very complex and difficult to trace from the original agent characteristics. Many slight variations on the genetic algorithm can add to system complexity, and most such additions are attempts to replicate biological processes or effects (Holland, 1995).

THIS PAGE INTENTIONALLY LEFT BLANK

III. COGNITIVE TASK ANALYSIS: TOOLS AND MODELS

A. MODELING EXPERT NAVIGATORS USING NDM

Previous researchers studied military navigators performing score orienteering tasks in a non-tactical environment (Peterson, Stine, and Darken, 2000). Military navigation is inherently different from civilian orienteering. It incorporates elements of small unit tactics, group leadership, and military mission planning that do not exist for the civilian orienteer. While the orienteer simply gets from point to point as quickly as possible, the tactical military navigator must conserve his patrol's energy for its mission, avoid detection by enemy forces, and maintain control over several soldiers. None of these vital considerations can be measured on a standard navigation course. This implies that, to properly evaluate navigators, the evaluator must observe them in the context of a tactical mission.

Naturalistic Decision Making theory fits tactical ground navigation because, by its definition, it relies on study of performance in the natural environment. Instead of attempting to isolate navigation as a task unto itself, NDM theory requires that the task be studied in the field in the context of a mission (Zsambok, 1997). Furthermore, the NDM framework defines properties of tasks and environments for which it is useful: dynamic and uncertain conditions; ill-defined, shifting, and competing sub-goals; continuous environmental feedback; high-tempo, high-stress, and high-stakes performance requirements, multiple players involved, and important organizational goals and norms (Orasanu & Connolly, 1993; Zsambok, 1997). These properties match the domain properties of tactical navigation quite closely (Peterson, Stine, and Darken, 2000).

A tactical patrol operates in enemy territory, moving through terrain with which it may be unfamiliar. The navigator is the member of the patrol primarily responsible for monitoring the environment as it changes and finding the correct route over terrain he may have never seen before. The route is carefully planned in great detail, but the plan may be changed without notice. Changes in the mission may cause the patrol to move to a different location, unexpected enemy activity may make the planned route unacceptable, or weather and vegetation may make some portions of the planned route impassable.

The patrol is always given a time to execute its mission, so obviously navigation to the objective must be completed before this execution time. Other sub-goals are not so well defined. The patrol must avoid enemy contact that would compromise its mission, but to move too far to avoid the enemy may take too long. Additionally, thick vegetation provides concealment for the patrol, but slows movement and saps strength from each patrol member that he may need on the objective. The navigator must monitor these and other factors and adjust his route accordingly.

Since the patrol is executing on a timeline, poor navigation can be catastrophic. Lost patrols may fail to get to the objective on time, compromising the mission. Navigators who cannot adjust to route deviations may wander into enemy positions or into a friendly sector where they may be mistaken for enemy patrols. Members of lost patrols quickly become frustrated and hostile to the navigator, and the resulting loss of morale may also adversely affect the mission.

B. DEFINING EXPERTISE

1. Exploratory Focus Group

In an attempt to clearly define the differences between novice, competent, and expert navigators, a focus group of ten instructors at the U.S. Army's 4th Ranger Training Battalion met on 30 SEP 99. These men had, at the time of the meeting, an average of 9 years of Army service and an average of 25 months as Ranger Instructors. Assumptions included that the navigators operate as part of a nine-man squad conducting a daylight tactical mission in wooded, rolling terrain similar to that found at Ft. Benning, GA. Furthermore, we divided the task into three distinct phases: planning, movement from Patrol Base to Objective Rally Point (ORP), and movement from ORP to objective. The result of the focus group is the identification of several performance characteristics that allow us to categorize a navigator as expert, competent, or novice.

2. Differences Between Experts, Competents, and Novices

a. Planning Phase

The group identified route selection, division of labor, and the quality of rehearsals as the three most important aspects of planning, if one assumes the perspective of evaluating proficiency in navigation. Route selection involves selecting a navigation technique - dead reckoning, terrain association, or some combination of the two. It includes selection of checkpoints, identifying specific terrain features to aid navigation, selecting boundary features, and adjustments for suspected enemy locations. Division of labor is the navigator's use of the other members of his unit to assist him, and how he uses the inputs they provide. Quality of rehearsals is a function of the quality of the

terrain model the unit constructs for the mission, and how the navigator uses it to convey information about the route to the members of the unit.

Novice navigators select straight-line routes and use azimuth and pace count to dead reckon their way from point to point. Routes are characteristically the shortest possible distance, regardless of terrain or the tactical situation. The novice tries to compartmentalize navigation as the first phase of his mission instead of something to be done throughout. Competent navigators incorporate terrain association, and their routes include checkpoints on identifiable terrain features and boundary features to indicate errors. Typically, the competent navigator uses roads as boundaries and man-made features as checkpoints. Additionally, they neglect to factor suspected enemy locations into their routes, as they also try to navigate first, then conduct the mission after navigation is done. Experts rely almost totally on terrain association. They use natural features as check points and boundaries, and structure their routes to avoid known or suspected enemy locations. The expert's route is typically longer than the novice's, but can be executed faster. The expert can deduce facts about the terrain that escape the competent, such as places the vegetation is likely to be thicker. By avoiding these places, his route again may be longer, but can be executed faster. Most importantly, the expert always limits his possible error by identifying natural boundary features along his route. His route thus becomes a corridor of movement, bounded by elevation changes, which allow him to detect deviation from the route. He plans to use the compass and pace count as guides and checks within the corridor, but only in a rough sense, relying on his skill in terrain association to navigate precisely. Most importantly, the expert does not try to decouple navigation from his mission. He realizes that navigation is woven into every

aspect of the mission and cannot be isolated, checked off a list of things to do, then forgotten.

Division of labor, in the novice case, turns into planning by committee, with each unit member's opinion weighted equally. The novice navigator asks for assistance from everyone, but does not know if the advice he gets is useful or counterproductive. Competent navigators tend to do everything themselves. They have enough experience to avoid novice mistakes, and realize that they can plan faster alone than in committee. Experts will evaluate each unit member's proficiency, then use good navigators to assist. They can quickly recognize bad advice, discount it, and assume added responsibility if needed. The expert has confidence in his own ability, and can quickly identify others whose input is trustworthy.

The quality of the unit's terrain model provides key clues to the proficiency of the navigator. A novice's terrain model is flat, emphasizes linear features like roads and creeks, and generally consists of a series of azimuths and distances connected by string or chalk lines. His rehearsal is a verification that each unit member has memorized each azimuth and distance. The competent navigator provides as much relief as time allows in his terrain model, and clearly indicates the check points his route includes. His rehearsals emphasize not only direction and distance but also terrain features that act as boundaries and checkpoints. The expert's terrain model is not significantly different from the competent's, but he uses it in a distinctly different manner. The expert's goal is to create a mental picture of the terrain in the mind of each unit member. He emphasizes the boundaries of the planned corridor, easily identifiable features to the left and right along the way, and areas where enemy activity may cause

changes. Experts use changes in elevation to explain the route. They make it clear when the unit will be moving up or down hill, and when they will cross ridges or valleys. An expert may refer to man-made features so the unit members know to expect them, but his error checking and confirmation always use natural features.

b. Patrol Base to ORP Phase

Key proficiency indicators during movement from the Patrol Base to the ORP include the flow of the mission, route changes, and error checking. Flow of the mission is determined by frequency of halts and actions at check points. Route changes include how often the planned route changes, why it changes, and how the navigator adjusts. Error checking includes not only how the navigator knows he's not where he wants to be, but how he confirms that he's right.

Novices stop the entire unit frequently and for long periods of time, consulting with other unit members each time. Fear and confusion is apparent on their faces. They then try to move faster to make up the time lost at halts. Competent navigators move their units at a maintainable pace, stopping at check points to conduct map checks. They confirm the checkpoints by terrain association in halts of no more than one to two minutes, then continue moving. A mission led by an expert navigator flows smoothly from Patrol Base to ORP. The expert already has the corridor memorized, so he doesn't need frequent map checks. He continues moving through checkpoints, checking his map as he walks. The expert keeps moving at a steady pace until planned halts.

The group estimated that during one of every three missions, something will happen to make the unit change course. The novice will not change at all unless

there is clearly no alternative. When he has no choice, he conducts a long halt, consults with the other unit members, and recalculates azimuths and distances for the adjusted route. The competent navigator will change his route any time the enemy situation or unexpected terrain (such as cleared woods or rushing intermittent streams) dictates. He will also conduct a halt, then plan a new route with new checkpoints to aid in terrain association. The expert will frequently change routes, not only in the aforementioned situations but any time an easier or better route presents itself. He is not afraid to move outside his planned corridor, and will do so after only brief halts to establish boundaries in his mind.

The only error-checking tool the novice uses is his pace count. He trusts his pace count, or another unit member's, before he trusts the map and his own skill in reading it. Novices perceive errors when they estimate distances different from those they planned. Competent navigators know they're in error when they miss expected checkpoints or cross boundary features. While they recognize errors far more quickly than novices, they cannot guess how far off they are. They rely almost totally on visual cues to confirm that they are correct, which causes them to lose confidence at night. The competent navigator is reduced to simple dead reckoning at night and is consequently often no better than the novice. Since they plan detailed movement corridors, experts have built error checking and confirmation into the route. Even when they are off, they know what the maximum error could be at any given point. The expert not only uses visual cues, he can feel the ground beneath him, knowing which foot should be at a higher elevation than the other based on his mental picture of the movement corridor.

Thus, when he loses the visual cues at night, he can still move his unit effectively, although much more slowly.

c. ORP to Objective Phase

Movement from the ORP to the objective includes recognition of the ORP location, conduct of the leaders' reconnaissance of the objective, and selection of release points and rally points. ORP recognition includes both where the ORP is placed and how the navigator selected the spot. Leaders' reconnaissance includes movement to the release point and confirmation of the objective. Selection of release and rally points involves what criteria the navigator uses to select them, and under what conditions they are useful.

The novice places his ORP at the exact spot his pace count coincides with the planned distance from the last checkpoint. He does not care if there is a better location nearby, but only that he can now transition from navigating to mission execution. The competent navigator will stop at the same spot, but look around for the best location to place an ORP. After a brief halt, he will move the unit to this location and establish his ORP. Then, he transitions to mission execution. The expert never performs such a transition, because he begins to look farther ahead as he enters the last leg of his corridor. He searches along the corridor for the best location, which may be farther away from or closer to the objective than planned depending on the terrain. He navigates to and selects an ORP that best suits his mission. To do this, he can't compartmentalize navigation and mission execution, but instead keeps mission concerns in mind throughout the movement.

Novices conducting leaders' reconnaissance move, by the most direct route possible, to the objective, quickly confirm it, and move directly back to the ORP without establishing a release point. They take less care in navigating, since they assume that they can navigate easily over the relatively short distance from the ORP to the objective and back. The competent navigator realizes that he must take more care, not less, and will ensure that he does not become disoriented on his way back to the ORP. He will carefully confirm that he is at his objective, then return. The expert again thinks several steps ahead. He too is very careful to navigate exactly. He will look for rally points and release points on the way to the objective, confirm the objective from multiple locations around it, then backtrack through all his release and rally points to ensure they are visible from both sides as he returns to the ORP.

Novices often disregard release and rally points, and put them in planned locations only (if they implement them at all). Thus, their rally points are not easily identifiable. Competent navigators take care to make their release and rally points identifiable, but forget to check if they are recognizable from multiple directions (especially coming and going), and that they can be seen at night. Experts designate release and rally points that can be easily identified from multiple directions, day and night.

d. Discussion

Three key abilities stand out as vital to military, mission-oriented navigation: perceiving elevation changes, integrating navigation into the mission as a whole, and creating a mental image of the terrain.

The ability to see and feel elevation changes was key to each phase of navigation, and stands out as the key terrain indicator by which experts navigate. Novices don't recognize the importance of elevation. Competent navigators can see and feel large changes. Experts are attuned to very small changes by sight and feel, and can use these changes to bound errors.

Novices and, to a lesser extent, competent navigators try to compartmentalize navigation as a task to be completed before mission execution begins. Experts realize that this cannot be done, and that navigation is a part of the mission as an indivisible whole, much like security.

Experts can look at a two-dimensional representation of terrain (a map) and create a mental picture of the three-dimensional real world they will encounter. Competent navigators can picture major terrain features, but their mental images lack the detail characteristic of experts. Novices lack this skill altogether.

3. Identification of Experts

Interviewing true experts in any domain is far easier said than done. It is often incredibly difficult to find actual experts, or to differentiate between them and the great mass of competent and proficient performers. The researcher must clearly define expertise and, even more important, clearly articulate specific criteria to identify individuals who meet the definition. Many studies involving military subjects fail to clearly define expertise. Experience alone is not a measure of expertise, nor is status as an instructor. In most cases, true experts in military skills are in units with real-world missions. Additionally, the best instructors are not necessarily the best performers. Since the very definition of expertise is intuitive recognition coupled with intuitive action, the

expert often finds it very difficult to describe his skills to novices or competents (Dreyfus, 1997). However, some programs definitely produce better performers than others, the focus group was asked where to find the best Army navigators.

The Ranger Instructors identified the Special Warfare Instruction Center at Camp Mackall, NC as a good source of experts. The average Special Forces Qualification Course student is more experienced in both small unit operations and land navigation than his counterpart at the Ranger Course. Those candidates with insufficient small unit experience or leadership ability are not admitted. Compared to the general population of Army navigators, Special Forces students are very experienced and generally more proficient. Additionally, they must pass a demanding orienteering-style course to continue training. Students identified by Special Forces instructors as among the top ten percent are therefore very likely to be true experts, since they would be in the top one percent of all military navigators.

C. KNOWLEDGE ELICITATION USING CDM

Initially, the intent was to directly observe the navigators while patrolling. In the natural patrol setting, especially in a school, it seemed that observation would be too obtrusive and organizationally disruptive. CDM seemed ideal for knowledge elicitation, since it provides enough flexibility for its adaptation to the specific needs of the researcher, while it also provides the structure necessary for getting the most out of each interview.

1. Planned vs. Actual Protocols

On 7 and 9 December 1999, a two-member team interviewed eight soldiers at the U.S. Army Special Forces Qualification Course, Phase I. All were identified as expert navigators by course instructors after approximately 21 days of training.

An instructor accompanied each student-led patrol. Immediately after the patrol finished its mission, the instructor would evaluate the leaders and critique the overall conduct. Depending upon the instructor's preference, we planned to interview the patrol navigator either during the instructor's critique or immediately following its conclusion.

Our team planned to begin the interview within an hour of each patrol's end and conduct it in 75 minutes according to the following protocol: 1. Meet with instructor to identify key difficulties, key decision points. 2. Elicitor orients participant to the patrol just completed as the patrol of interest. 3. The participant recounts the entire patrol. 4. Elicitor retells the story back to the PL. This allows both PL and interviewer to arrive at a common understanding of the sequence. 5. Elicitor and participant build a time line of the sequence of events. The timeline will include decision points, inputs to each decision point and actions taken as a result of each decision. 6. Elicitor asks probe questions to deepen his understanding of the navigation.

Ideally, we would like to have met with the instructor immediately before we conducted the actual student interview; we hoped that the instructor's comments would help us focus the interview on the key decision areas of that particular patrol episode. In practice, this was not practical and we were unable to meet with the instructors. The window between the end of the patrol and the beginning of our interviews was quite

narrow and the instructors were busy enough preparing themselves for their own critiques.

We intended to generate two artifacts during each interview. The first was the participant's sketch of the patrol; the second was a timeline, with key decision points indicated on it. After attempts to produce both artifacts, we dropped the timeline and focused effort on the patrol sketch. It seems that the sketch afforded a focal object for the discussion, while the timeline served to scatter the discussion too much.

2. Setting and Conduct of Interviews

Most of the interviews were conducted indoors. Those conducted outdoors enjoyed clear, dry and warm weather. The field interviews were quite different from those done indoors. Indoors, we had two butcher pads with easels and lighting. In the field, we lay in the dirt, shining a flashlight on the student's map while scribbling notes on a clipboard. This method had some advantages in that it focused and sped the interview; unfortunately, we did not get as much detail nor did we obtain patrol sketches.

THIS PAGE INTENTIONALLY LEFT BLANK

IV. REPRESENTATION OF EXPERTISE

A. GENERAL RESULTS

1. General

Complete description of expertise requires four separate representations. First, a high level overview explains the important characteristics. Second, a cue inventory shows the environmental factors experts process to make decisions. Third, specific instances from the interviews illustrate how experts apply each of the three variations of the RPD model. Finally, a Key Decisions Requirements Table integrates the three functions into one model of decision making.

Although expertise is a highly individual phenomenon, and is developed differently by each individual, all the experts interviewed had four basic characteristics in common: they rely on high-fidelity mental maps; they blend multiple cues; they adjust and recalibrate tools dynamically; and they visualize spatial information.

2. Mental Map Fidelity

During mission planning, the expert spends a great deal of time and effort creating a highly detailed, three-dimensional mental map. This is much more than simple memorization of the paper map, as the expert visualizes the terrain as it will appear, including features hinted at but not included on the paper map. The mental map includes vegetation, relative terrain elevation, roads by type and quality, streams and lakes, and any man-made structures. During the execution phase, experts refer to the paper map only in extreme cases, and trust the mental map completely.

The mental map focuses on the planned route, and builds from there. The route includes a general compass azimuth and measured distance, and indeed this is exactly

how the expert describes it to other patrol members. However, the expert views the route more as a corridor than as a line on the map, with key terrain features on each side serving as lateral boundaries. All experts included easily recognizable checkpoints to separate route legs, and emphasize these to other patrol members. Routes always incorporate all key terrain features in visual range, including roads and streams, man-made objects, and changes in elevation.

To describe the route to the other patrol members, the navigator creates a three-dimensional terrain model. Given sufficient preparation time, the quality of the terrain model seems to be closely related to the quality of the mental map, and as such is a clear indicator of proficiency in navigation. This corresponds with the Ranger Instructors' assessment, and suggests that the act of building the model helps reinforce the map in the navigator's mind.

The expert's mental map includes far more information than just the route. It includes detailed information on the entire area of operations. This level of detail allows the expert to make dynamic changes to the route without consulting the paper map. Furthermore, experts can use it to adjust navigation tools, mentally simulate the consequences of decisions, and generate stories to explain expectancy violations during the execution phase.

3. Blending of Cues

While walking, experts process information from the environment and compare it to the mental map. They use three major and two minor cues, and assign relative weights to each based primarily on environmental conditions. However, it seems that different experts may use significantly different relative weights, even in the same environment.

Personal preferences, training, experience, or some other factors may cause these differences, which may be noted with further research.

The major cues are terrain features, compass azimuth, and pace count. The expert always monitors all three. However, his degree of reliance on each one is based on weather, vegetation, light, and visibility conditions. For example, in daylight with little vegetation, the expert may use terrain features more than the other two, since he can see all features. However, at night, or in thick fog, he may not be able to see the useful terrain features and must rely more on azimuth and pace count.

The minor cues are tactical and mission considerations. These affect the expert in more subtle ways, but are still considered continuously. For example, proximity to a suspected enemy location may cause different action, as will contact with enemy forces. Also, difficult routes may tire the patrol more quickly than expected, so the navigator must compensate to preserve energy for the actions on the objective.

4. Dynamic Adjustment

Experts can dynamically calibrate and correct navigation tools. They keep a pace count to estimate distance traveled and frequently check a magnetic compass heading. If either of these provides information in conflict with the mental map, the expert can approximate the error and recalibrate the tool on the fly. Experts frequently measure, over a known distance (usually 100 meters), how many steps they take to cover that distance. However, pace counts may vary widely due to fatigue, visibility, and rough terrain. Experienced navigators can factor these into the pace count. Experts process pace count information, compare measured distance traveled with mental map distance, and adjust the pace count to correct any discrepancy. As the navigator walks, he cannot

follow a straight azimuth, but must move around trees, lakes, boulders, and other obstacles. Moreover, unexpected enemy contact may require the patrol to deviate from the planned route. In these cases, experts can mentally compute new azimuth headings and implement them without stopping. Experts view a halt for a paper map check as an abject failure. Frequent stops lower the patrol's confidence in the navigator and thus overall morale. Dynamic adjustment allows experts to minimize these stops. While the organizational expectation is minimal map checking, experts hold themselves to a higher standard - zero paper map checks.

5. Spatial Visualization

From map study alone, experts can visualize three-dimensional terrain. They can also, while walking, visualize how real terrain would look on a two-dimensional paper map. These two related skills are vitally important hallmarks of good navigators. The first, known as map to ground, is primarily important in the planning phase, as it allows the expert to create his detailed mental map. It enables the navigator to select the proper route and create a useful terrain model to explain the route to the rest of the patrol. Conversely, experts use the second skill, ground to map, during the execution phase. It allows comparison of real terrain to mental map, and allows the expert to make necessary azimuth adjustments dynamically. This is a continuous process, and is beyond the capabilities of novice navigators. Experts mention the development of spatial visualization as a key element in the development of expertise, since without it the other skills cannot develop fully.

B. CUE INVENTORY

1. General

Adapted from Hoffman, et al, (1995) a cue inventory in Figure 6 shows cues grouped by category. The description of each cue listed in Figure 6 follows.

Category of Cue	Cue
Navigation Tools	Compass Azimuth
	Pace count
	Paper Map
	Mental Map
Environmental Conditions	Ground Slope
	Vegetation
Mission Conditions	Time
	Input from Other Patrol Members
Terrain Features	Road
	Body of Water
	Topography
	Man-made Feature

Figure 6. Cue Inventory grouped by Category

2. Navigation Tools

The navigator always carries a compass that is used to determine the direction of travel and direction to landmarks. Depending upon the task organization of the patrol, other members may be responsible for ensuring that the patrol travels on the intended azimuth. Expert navigators have a high tolerance for deviations from the intended azimuth since they combine the compass azimuth with other cues to maintain orientation. The navigator counts his steps and accumulates them while walking. As with azimuth, other patrol members may be tasked to keep a pace count for the navigator. Over practice, the navigator knows how many paces he must walk to cover one hundred meters. Since the stride length and rate vary based upon terrain, vegetation, slope, fatigue, speed of travel, visibility and weather, the pace count can be a misleading,

inaccurate cue. Expert navigators are able to recalculate their pace count dynamically as conditions change, giving them a more accurate, reliable count. The map is the primary source of information used during route planning, and it provides direct vegetation and elevation data. The navigator usually carries a map in his pocket while walking. Experts report that while walking, they make map checks extremely rarely, only as a last resort.

During route planning, the navigator studies the paper map and generates a corresponding mental map. The specific symbols, format and contents of the mental map are unclear. While walking, most navigators compare their physical surroundings to their mental map, so it is the main source of information once the patrol begins movement. Experts are able to rapidly generate detailed mental maps, hence relieving their dependence on the paper map.

3. Environmental Conditions

Through map study, experts are able to visualize the terrain as if they were walking over it. One component of this visualization is the slope of the ground. Not only is the slope of the surrounding terrain important, but also the slope of the ground immediately underneath the navigator's boots. While walking, experts are highly sensitive to changes in the ground slope, and they use both visual and kinesthetic cues to monitor the state.

The paper map usually depicts vegetation as green areas, but the map does not characterize the type of vegetation. From map study, the navigator will generate expectancies of how the vegetation will vary across the proposed route. These expectancies are critical as the type of vegetation has strong influence on cover,

concealment and movement effort. Experts are able to make fine discriminations in vegetation quality.

4. Mission Conditions

During route planning, the navigator carefully considers the time constraints dictated by the mission. The navigator must select and then execute a route that will result in the patrol's arrival at the objective at the prescribed time. An ability to estimate travel times from map study enables accurate planning. Route descriptions included spatial and temporal components. While walking, the navigator's expectancies about the order of landmarks and terrain features all are time-stamped. While distance traveled, environmental conditions and fatigue level can give clues about the time, a watch is the primary source.

While the navigator is primarily responsible for route planning and execution, other patrol members are also involved. The patrol leader is the main source of mission-related information. At any time, he can command that the navigator increase or decrease the pace of movement or that the route be changed. In addition, the leader can assign navigation-related duties, such as keeping the pace count or azimuth, to other patrol members.

5. Terrain Features

Navigators commonly use roads as landmarks and checkpoints, and maps usually depict roads. Depending upon the environment, the roads may be paved, gravel or dirt. For security reasons, navigators strongly avoid improved or paved roads. Patrols cross roads; they never travel on a road as civilians typically would. Experts are able to make fine discriminations between types of roads, such as paved, improved, main or secondary.

One reason why roads can be tricky landmarks is that often a patrol will encounter a road that does not appear on the map. These are commonly referred to as "false roads" because the road they see is not the planned landmark. Experts are able to process cues such as the road's bends, taper, slope and evidence of traffic to correctly identify false roads.

Water features such as river, creeks and streams are commonly used landmarks. At times larger bodies of water, such as ponds and lakes are used. The level of water is an important cue that skilled navigators report using. The extreme case of a dry creek is common during the dry seasons. As with roads, navigators contend with the possible presence of false water bodies. Recent rainfall can channel water through low ground that is not depicted as a water body on the map; conversely, dry environmental conditions can leave the low area depicted on the map as water, completely dry. Experts are able to make use of a combination of cues to determine which dry creeks are false and which are true.

The terrain's relief and elevation are the most commonly used navigational features. As the scale of the map increases, there is less detail depicted by the map, so only the most prominent hills can be pinpointed by map study. The navigational skill of reading a two-dimensional map and visualizing a three-dimensional space hinges on the ability to interpret map relief and elevation symbology. Experts can do it. Again, navigators must contend with possible false hills, valleys, draws and ridges. While walking, experts are able to pick out the true features and discount the false ones. They know which features are of sufficient size to appear on the maps they use, and can infer the smaller ones from the information contained on the paper map. They can compare

these inferred features with the terrain as well, and further distinguish false terrain features from relevant ones.

With the exception of roads, skilled navigators choose to attend to natural terrain features rather than man-made objects, such as buildings. This tendency seems to be rooted in the reality that combat operations and non-combat inhabitation can drastically affect the presence and appearance of man-made structures.

C. SITUATION ASSESSMENT RECORD

1. General

As described by Hoffman, et al. (1995), a situation assessment record highlights the points where the expert made a decision based upon a revised assessment of the situation. After examining the example presented there and comparing it to the elements of the RPD model as diagrammed by Klein (1998), it seemed that the RPD pattern could be used to describe expert navigation and describe the situation assessment record. There are three different variations of the flow through the RPD model, and each variation is related to the decision-maker's recognition of the situation. Variation 1 describes episodes where the expert recognizes a typical situation. The fact that the situation is typical means the expert takes action immediately, without thinking; the recognition of the situation primes the appropriate action. Variation 1 typifies the quick and accurate behavior frequently associated with expertise (Dreyfus, 1997).

Sometimes, even experts are faced with situations that are not immediately recognized as being typical. Here begins Variation 2. During these episodes, the expert directs mental effort to the process of recognizing the specific cues and patterns that comprise a situation. As the RPD model asserts, this diagnosis involves two mental

processes. First, the expert identifies the features of the situation; then he compares the present features to other situations to match the feature contents and arrive at situation recognition. Second, the expert creates a story to explain how the features might fit together and what actions might have caused the situation to arise. Often, the decision-maker will alternate between the feature matching and story creation processes, until the current situation can be categorized as being typical. Once the situation is recognized, then the expert takes action as in Variation 1.

Variation 3 begins with the recognition of a situation. However, unlike Variations 1 and 2, in these cases, the expert does not immediately know what to do. Mental effort is expended not on situation recognition but response evaluation. In some ways, the expert behaves as a competent performer would (Dreyfus, 1997). He must figure out what to do. The RPD model specifies that this action evaluation happens in a way that differs from traditional decision-making theory. Rather than simultaneously comparing multiple responses, the expert considers them singularly. The decision-maker mentally simulates forward from the current situation to the simulated outcome of the first action that comes to mind. If the outcome is workable, then the expert implements it. If not, then he discards it and considers another option. Sometimes, this simulation will identify an outcome that satisfies most of the relevant goals but not all of them. In these cases, the expert may make slight changes to the action and then rerun the simulation.

Obviously, the expert's recognition of the situation is the key. This recognition generates four "by-products" (Klein, 1998) or types of mental constructs useful to the expert's future performance: expectancies, relevant cues, plausible goals and typical actions. For a given situation, there are associated expectancies about what will happen

next. Sometimes, these expectancies are expressed in terms are relevant cues. The expert attends to the relevant cues to confirm or disconfirm the expectancies of the situation. The violation of an expectancy often triggers a new situation assessment. Also, the situation defines which goals are plausible. Decision-maker attention to relevant cues and input from the organization can cause the relative importance of these goals to shift, and sometimes these shifts will generate a new situational assessment. Finally, to achieve the goals, the expert has a set of typical actions associated with each situation. In Variations 1 and 2, the expert implements one action from this set without evaluating each possibility (Klein, 1998).

Specific items for each of the four by-products are illustrated in Figure 7. Next, one story, drawn from our interview data, illustrates each of the three variations.

The Standard By-Products of the Expert Navigator's Situation Assessment	
Expectancies Generated by evaluation of the situation with regard to the mental map.	**Relevant Cues** Selected from the cue inventory categories: • Navigation Tools • Environmental Conditions • Mission Conditions • Terrain Features
Plausible Goals Selected from the list of standard goals: • Maximize Speed • Maximize Stealth • Minimize Exertion • Maintain Orientation	**Typical Actions** The Standard Typical Action is one of Three Methods: • *Arrive at Checkpoint* Method 1. Confirm checkpoint if needed. 2. Reset pace count. 3. Change azimuth if needed. • *Confirm Route* Method 1. Maintain pace count. 2. Maintain azimuth. • *Error Recovery* Method 1. Confirm checkpoint if needed. 2. Reset tools if needed. 3. Map Check if needed.

Figure 7. Specific By-products of Land Navigation Situation Assessment from (Klein, 1998)

2. Situation Assessment Record, Variation 1

All of our participants operated under Variation 1 conditions most of the time. They recognized the navigation situation as being typical, and they just continued to navigate – walking and scanning the environment. The record is shown in Figure 8. The relevant cues, plausible goals and typical actions are all drawn from the standard sets, as listed in Figure 7. The story begins as the patrol moves from its starting point through checkpoint one and onto checkpoint two. The navigator initially expects to walk uphill, and then cross a road. As these expectations are met, the navigator is acting according to

the "confirm route" and "arrive at checkpoint" methods. After crossing the road, he navigates to checkpoint number two, consulting his mental map to update his expectancies for the respective leg of the route.

Example of Variation 1 – "I know the situation, therefore I know the course of action."
"Continue Mission"

Situation One: On course between start point and checkpoint 1
Relevant Cues: Standard.
Plausible Goals: Standard.
Typical Actions: Standard.
Expectations: Expect to begin by moving uphill, then we will cross a road.
Course of Action: Arrive at Checkpoint Method.

Situation Two: On course between checkpoint 1 and checkpoint 2
Relevant Cues: Standard.
Plausible Goals: Standard.
Typical Actions: Standard.
Expectations: After crossing the road, we will hit a draw. We will box around the draw. Then we should cross a road.
Course of Action: Arrive at Checkpoint Method.

Figure 8. Situation Assessment Record for RPD Variation 1

3. Situation Assessment Record, Variation 2

As shown in Figure 9, the example of Variation 2 comes from a participant who was able to recognize his own error and correct it dynamically, on the move without disrupting the flow of the patrol's movement; it is likely that the other patrol members were not even aware that the error occurred. The patrol was moving smoothly from checkpoint to checkpoint. Enroute to the patrol's sixth checkpoint, the navigator expected to cross two draws and then an improved road. However, an anomaly violated this expectation, as he crossed a secondary road after moving only 100 meters past the fifth checkpoint. After matching the relevant features of the situation, he considered a story in which they had crossed the improved road too far north. This was caused by

mistakenly cutting the last leg short before changing heading. The source of the error was that his pace count had become mis-calibrated, likely due to fatigue. He verified this story against his mental map and the visible terrain features, and assessed the situation to be error recovery. Relying on his detailed mental map, he knew where the patrol was but his pace count was off. Remembering from his mental map, the distance between the improved and secondary roads, he re-calibrated his pace count and later confirmed his revision.

Example of Variation 2 – "What is the situation?"
"I'll just do a quick dynamic pace count recalculation..."

Situation One: On course between checkpoint 2 and checkpoint 3
Relevant Cues: Standard.
Plausible Goals: Standard.
Typical Actions: Standard.
Expectations: We will cross another road. The Pace count here should be 300m. We will next be able to see a hill. Next, we should cross a major road.
Course of Action: *Arrive at Checkpoint* Method.

Situation Two: On course between checkpoint 3 and checkpoint 4
Relevant Cues: Standard.
Plausible Goals: Standard.
Typical Actions: Standard.
Expectations: We will identify a bend in the road at 450m.
Course of Action: *Arrive at Checkpoint* Method.

Situation Three: On course between checkpoint 4 and checkpoint 5
Relevant Cues: Standard.
Plausible Goals: Standard.
Typical Actions: Standard.
Expectations: We will cross a major road at 1000m.
Course of Action: *Arrive at Checkpoint* Method.

Situation Four: On course between checkpoint 5 and checkpoint 6
Relevant Cues: Standard.
Plausible Goals: Standard.
Typical Actions: Standard.
Expectations: We will cross two draws. Then we will cross an improved road.
Anomaly: We crossed a secondary road at 100m.
Diagnose:
Feature Matching: Traveling on the right compass heading. Pace count is 100m. Crossed road. Have not crossed draws. (Matches these features to the features of his mental map.)
Story: On our last leg, we must have stopped too far north. Then on this leg, we crossed the major road too far to the north/east.

Situation Five: Off course between checkpoint 5 and checkpoint 6
Goal: Reorient and compensate for error.
Course of Action: *Error Recovery* Method: Recalculate pace count, mentally change the route.

Situation Six: On course between checkpoint 5 and checkpoint 6
Relevant Cues: Standard.
Plausible Goals: Standard.
Typical Actions: Standard.
Expectations: We will cross two draws. Then we will cross an improved road.
Course of Action: *Confirm Route* Method

Figure 9. Situation Assessment Record for RPD Variation 2

4. Situation Assessment Record, Variation 3

This example of Variation 3 begins with an anomaly, as presented in Figure 10. Enroute from checkpoint four to checkpoint five, the patrol walked 400 meters and the navigator expected to arrive at checkpoint five. He did not. From feature matching and story generation, the navigator accurately reassessed the situation and realized that the patrol had not reached the planned checkpoint four. However, the solution was not immediately obvious, since checkpoint five was in enemy territory. The patrol leader had established a temporary defensive position at checkpoint four and currently, most of the patrol members were located there, preparing to engage the enemy. Although the navigator recognized the situation, and he knew where they were, the situation was not typical and he immediately evaluated possible actions. He first considered going back to checkpoint 4 and relocating it, but then he discounted it since it would cause too much confusion and lower morale, which could jeopardize the patrol. Relying on his detailed mental map, he mentally constructed a new route from the present checkpoint four to the desired checkpoint five and realized that would be the simplest action. But, what if he got shot? The rest of the patrol would not know where they were. So, he decided to change the route and inform the patrol leader of the change; the leader would then decide how to disseminate the route change.

Example of Variation 3 – "I know the situation…what do I do about it?"
"Hey, we're too far from the checkpoint…"

Situation One: On course between checkpoint 2 and checkpoint 3
Relevant Cues: Standard.
Plausible Goals: Standard.
Typical Actions: Standard.
Expectations: Should enter triangular open area between roads. Should then cross major road.
Course of Action: Confirm Route Method.

Situation Two: On course between checkpoint 4 and checkpoint 5
Relevant Cues: Standard.
Plausible Goals: Standard.
Typical Actions: Standard.
Expectations: Checkpoint 5 should be 400m away on set azimuth.
Anomaly: At 400m, did not hit checkpoint 5.
Diagnose:
Feature Matching: Two hills, one to our left and one to our right. Our compass heading is correct. Pace count is 400m. Estimate that checkpoint 5 is still 300m away distant.
Story: We must have misplaced checkpoint 4. On our last leg, we did not go far enough east. That means we are 400m west of our planned location.

Situation Three: Erroneous checkpoint 4
Goal: Maximize Stealth. Minimize patrol confusion.
Typical Actions: None.
Evaluate Actions: Go back and move patrol's defensive position (Mental Simulation).
Will it work: No… patrol is preparing to engage the enemy.
Evaluate Actions: Change route from checkpoint 4 to checkpoint 5 (Mental Simulation).
Will it work: Yes, but…Patrol Leader must be informed.
Evaluate Actions: Change route and inform Patrol Leader. (Mental Simulation).
Will it work: Yes.
Course of Action: Change route and inform Patrol Leader.

Situation Four: On course between checkpoint 4 and checkpoint 5
Relevant Cues: Standard.
Plausible Goals: Standard.
Typical Actions: Standard.
Expectations: Checkpoint 5 should be 700m away on new azimuth.
Course of Action: Confirm Route Method

Figure 10. Situation Assessment Record for RPD Variation 3

D. KEY DECISION REQUIREMENTS

A listing of the key decision requirements is useful to determine which decisions are particularly critical, and which skills enable the decisions (Hoffman, et al., 1995). Such mapping helps direct training resources. Figure 11 presents a portion of the key decision requirements for land navigation. For each decision, we propose an explanation

for why the decision is difficult and how experts make it. The final column traces the process back to the four key mental processes introduced earlier.

What Is The Decision?	**Why Is It Difficult?**	**How is it made?**	**What enables decision?**
Selecting the Route	Maps do not provide enough information; difficult to estimate environmental conditions	Mental Simulation	High-fidelity mental map; spatial visualization
Memorizing the Route	Limited Time; must decide which features to memorize	Mental Simulation	High-fidelity mental map; spatial visualization
Identifying Terrain Feature	Environmental conditions can limit visibility	Compare environmental cues to expectancies	High-fidelity mental map; cue blending; spatial visualization
Discounting False Terrain Feature	Maps do not show all possible terrain features; must rely on multiple cues; requires multi-tasking attention	Compare environmental cues to expectancies; Mental Simulation	High-fidelity mental map; cue blending; recalibrate tools; spatial visualization
Recognizing misorientation	Difficult to recalibrate tools enroute; difficult to identify terrain features	Compare environmental cues to expectancies; Mental Simulation	High-fidelity mental map; cue blending; recalibrate tools; spatial visualization

Figure 11. Key Decision Requirements from (Hoffman, et al., 1995)

V. ROUTE PLANNING EXECUTABLE MODEL

A. AGENT-BASED ROUTE PLANNER

1. Focus on Route Planning

An executable model of the complete representation of expertise, including mission planning, route selection, rehearsals and briefings, sensory perceptions, and route execution, would be a very large and complicated project. Moreover, some elements of the RPD model are extremely difficult to translate to computer code. For instance, the story generation function requires creative and original thought from the expert, something no computer currently in existence can do. Even if the uniquely human thought processes could somehow be replicated, the physical environment in its infinite complexity could not. This means that the computer constructs perceptions would not match those of human experts in real environments. While approximate sensory perceptions are good enough for most simulations, a significant part of expertise is the ability to perceive subtle nuances in the environment - exactly the things that are not included in approximations.

Focus on the route planning portion of expert performance ameliorates these problems. The expert plans his routes on a 1:50,000 scale map, usually without being able to physically walk the ground over which he will travel. The map is a very rough approximation of the real world, one that allows the novice to perceive the same colors and symbols as the expert. Even so, the expert plans routes that differ significantly from those planned by novice or competent navigators. Planning does not explicitly require

story generation or subtle perceptions, and so is easier to approximate with machine logic. For these reasons, the executable model is solely a route planner.

2. Selection of an Agent-based Representation

Most expert systems use rule based artificial techniques, or derivatives of these such as case based reasoning. These approaches require huge databases of rules, covering all possible combinations of environmental factors. The rule base is not valid outside the environments for which it was constructed, as there are no rules to cover new environments. Agent-based systems are different in that they do not require rule bases for each environment, only decision vectors, which control action, based on inputs from the environment. The decision vector need not be changed by the programmer, but can be changed by the genetic algorithm in each new environment. Agents can therefore operate in any environment. Given sufficient repetitions, and a scoring system that correctly rewards desired performance, they will become experts. The exploratory focus group supplied a detailed description of expert route planning, which forms the basis for a decision vector in the agent model and the scoring system controlling the genetic algorithm.

B. TERRAIN REPRESENTATION

1. Organization of OpenFlight Files

One of the most common formats for models of real terrain is the Multi-Gen/Paradigm, Inc. OpenFlight™ format, a binary file format that describes geometry, colors, and textures. OpenFlight files have a hierarchical structure and individual polygons in each model can be grouped together at the developer's discretion. Multi-

Gen's Creator package, used to create OpenFlight files, will export OpenFlight to VRML97, among other formats.

A route planning agent needs to operate on a terrain database that would give the agent roughly the same information that humans get from a 1:50,000 map. Digital Terrain Elevation Data, Level 2 (DTED-2), from the National Imagery and Mapping Agency provides elevation data at this level of detail (NIMA, 1999), so the model was created with this data. This model is a desert database, so wooded areas, paved and dirt roads, and streams were added using Creator. The resulting model is organized into six separate terrain types: open areas, thick woods, light woods, paved roads, dirt roads, and water. Coupled with the elevation replicated in the geometry, the model gives a reasonable, but still somewhat simplified, version of a military map.

The organization of the database allows the agent to query it and get access to the type of terrain at any point, just as a human navigator can read the colors on the map. In fact, this is exactly how it is done. Since color is an easily accessible field in a VRML file, and the VRML file structure mirrors the OpenFlight file from which it was created, the agent simply asks the database for the color of the polygon at a requested location. Based on the color, the agent then accesses terrain attributes in the underlying Java classes for the appropriate terrain type.

2. User Interface

The user can control the development of expertise through the scoring system, as he sets the priorities for the agent navigator. He specifies the relative importance of avoiding enemy contact, fastest route, and least difficult route. By default, all three are set to equal importance. Agents whose decision vectors produce routes in keeping with

the user's priorities will score higher than other agents, and thus will reproduce with greater frequency.

Additionally, the user sets the maximum number of iterations in the "warm-up" period, during which the genetic algorithm modifies agent structure, and the known and suspected enemy locations.

3. Enemy Locations

The user sets the enemy locations and sensor ranges, and can change them at any time between iterations. Agent navigators will use enemy sensor ranges when planning routes as the minimum distance to keep between themselves and the actual location of the enemy element. Enemy infantry has a 500-meter sensor range, and mechanized/armor has a 3000-meter sensor range.

C. AGENT DESIGN AND STRUCTURE

1. Decision Vector

Agents in this model are based on the ISAAC agents developed by Ilachinski (1997). Each agent moves in five meter steps, in one of eight directions - the four cardinal directions and the four in-betweens. They contain a decision vector, which is a set of double precision, floating point values consisting of six elements. Each of these elements assigns a weight to the agent's propensity to take a certain action. Each agent has a different decision vector, and this is what makes the agents act differently, and makes the genetic algorithm effective. Just as biological systems depend on genetic diversity for survival, the greater the variation in decision vectors, the more adaptation is possible. Individual agents are controlled and scored by a MoverManager object, a software construct that also administers the genetic algorithm.

Decision vector values are in effect the relative importance the agent places on its different, and often competing, goals. Larger values for specific elements result in a greater propensity for the agent to satisfy the corresponding goal to the exclusion of the others. These goals include minimizing time of movement, minimizing distance traveled, minimizing total cumulative change in elevation, maximizing cover and concealment, avoiding enemy contact, and minimizing linear danger area (road and stream) crossings. The decision vector elements are:

ω_1: move to goal (shortest distance)

ω_2: move to least elevation change

ω_3: move to avoid enemy

ω_4: move along the fastest route

ω_5: move to cover/concealment

ω_6: move to avoid road/stream crossing

These elements came directly from the focus group's answers to the question "How do you tell experts from novices during the planning phase?", confirmed during the interviews with actual performers.

2. Perception Vector

The perception vector is the set of values that describe the important attributes of a specific location in the environment. "Important", in this case, is defined as affecting the agent's ability to meet one or more of its goals. Obviously, the perception vector changes based on the agent's location in the environment. It is calculated for each pairing of current and proposed locations. Thus, since the agent can move in eight possible directions with each step, each agent calculates eight perception vectors for each turn.

Perception vectors do not depend on the decision vector or any other agent attributes, so two agents occupying the same location would have eight identical perception vectors (if they were heading to the same goal). The perception vector is:

α_1: difference in distances to goal between current and proposed locations

α_2: change in elevation from current to proposed locations

α_3: 0 if not in enemy sensor range, else difference in distances to known enemies

α_4: time to move from current to proposed

α_5: change in amount of cover from current to proposed locations

α_6: 0 if proposed location is not a road or stream, penalty value if it is

3. Establishing Checkpoints

The agent's goal is not always his final destination, but shifts depending on his location in the environment. The goal at any given time is the next element in a set of checkpoints, defined before the agent begins planning the route. When the program starts, the MoverManager creates a set of possible checkpoints, a two-dimensional array of locations 1000 meters apart that effectively form a rectangular grid over the entire map. Each checkpoint stores its own terrain type and elevation for perception vector calculation, and a gradient value for route optimization.

Each agent calculates its set of checkpoints using the decision vector/perception vector method. The agent starts with his desired finish point, then works backwards to the start point, defining checkpoints every 1000 meters along the way. Routes shorter than 1000 meters contain no checkpoints, just the start and end points. The agent subtracts the scalar product of the decision and perception vectors from the gradient value, then sends the gradient and checkpoint set to each of its adjoining checkpoints.

Each checkpoint compares the gradient value received to its gradient value stored. If the received value is higher, the checkpoint adds itself to the set of checkpoints, stores the received value, then forwards the gradient and checkpoint set to each of its neighbors. If, on the other hand, the received value is lower, the checkpoint takes no action.

When a gradient reaches the agent's start location, that location compares its value to the highest gradient value received to date. The highest value received indicates the best set of checkpoints. By default, the start point stops the calculations after it receives the 50^{th} gradient, but the user can adjust this value if desired. This method is designed to prevent the "box canyon" effect of following a temporarily advantageous path to a very poor position.

4. Movement Using Lowest Penalty Function

Agents move five meters at a time, in the one of the eight directions corresponding to the lowest value of a penalty function. The penalty function is the scalar product of the decision and perception vectors, with the goal being the next checkpoint in the set of checkpoints. Since each agent has its own set of checkpoints, calculated using its decision vector, agents may have different checkpoints, and thus different goals. When the agent moves to within 200 meters of the checkpoint, the goal for the next turn becomes the next checkpoint in the set. This allows the agent to "round corners" if the route is more advantageous. Figure 12 illustrates a sample penalty function calculation. Lightest gray areas are lightly wooded terrain, dark gray represents heavy woods, and mid-intensity gray represents open areas. Assume that the agent's personality vector is as follows:

$\omega_1 = 0.75$

$\omega_2 = 0.30$

$\omega_3 = 0.50$

$\omega_4 = 0.25$

$\omega_5 = 0.67$

$\omega_6 = 0.45$

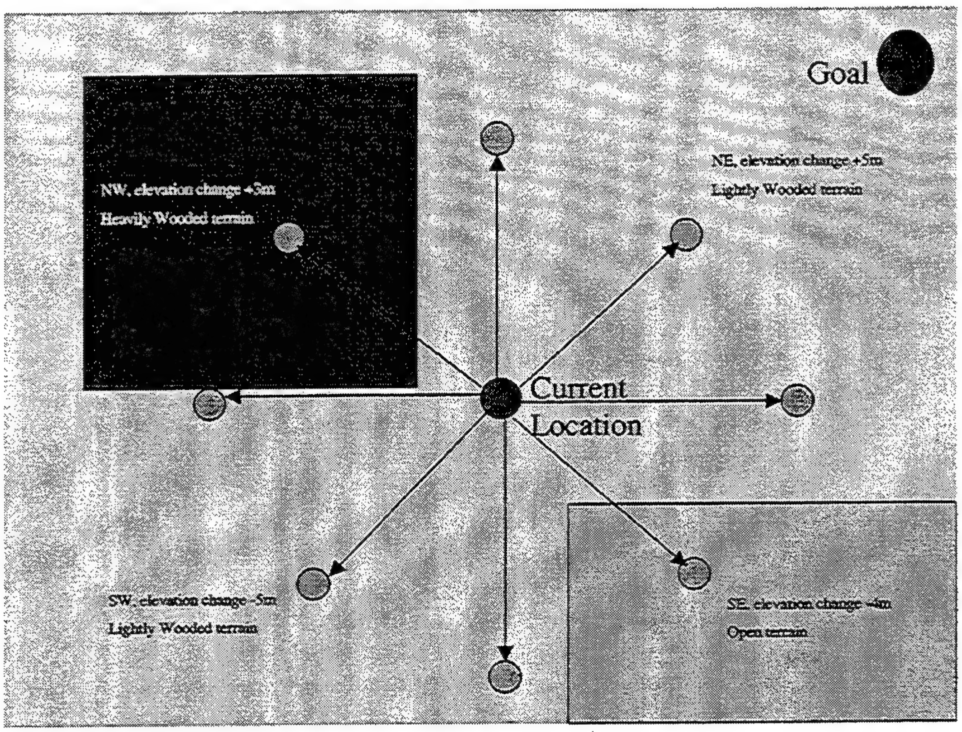

Figure 12. Sample Penalty Function Calculation

For the possible movement location NE, the perception vector is:

$\alpha_1 = 5m - 10m = -5.0$

$\alpha_2 = 5m = 5.0$ (sign does not matter - this calculation uses absolute value)

$\alpha_3 = 0$ (there are no enemies in the example)

$\alpha_4 = 7.5sec * 1.25$ (lightly wooded movement factor) $= 9.375$

$\alpha_5 = 0$

$\alpha_6 = 0$

For the location SE, the perception vector is:

$\alpha_1 = 12m - 10m = 2.0$

$\alpha_2 = 4m = 4.0$

$\alpha_3 = 0$

$\alpha_4 = 7.5sec * 1.0$ (open terrain movement factor)

$\alpha_5 = 1.0$

$\alpha_6 = 0$

Thus, the penalty function value for location NE is:

$$(0.75)(-5.0) + (0.3)(5.0) + (0.5)(0) + (0.25)(9.375) + (0.67)(0) + (0.45)(0) = 0.09$$

The corresponding value for location SE is:

$$(0.75)(2.0) + (0.3)(4.0) + (0.5)(0) + (0.25)(7.5) + (0.67)(1) + (0.45)(0) = 5.25$$

Therefore, the agent will move to location NE over location SE.

D. SCORING AND GENETIC ALGORITHM

1. Multiple Agents and Scoring

Agents must be scored over many routes, then ranked from highest scoring to lowest, to identify the most and least successful navigators for action by the genetic algorithm. The system creates twenty agents and randomly generates start and end points for the first route at startup. All twenty agents start and finish at the same points for each iteration. After all agents complete a route, the MoverManager randomly generates new start and finish points.

As each agent moves through its route, it keeps a running total of its distance traveled, movement time, and total change in elevation. The MoverManager maintains a security score for each agent based on the types of terrain its route crosses. Steps in open areas add ten points to the security score, while crossing roads or streams add twenty. The higher the security score, the less the agent moves securely.

When all agents reach the finish point, the MoverManager calculates the mean distance traveled, mean time to complete the route, mean security score, and mean elevation change, plus the standard deviation of each of these averages. Each element of the score is the number of standard deviations from the mean of all 25 agent scores, and as such can be positive or negative. The MoverManager calculates how far each element of each score is from the mean for that element. Then, it computes the adjusted score by multiplying the different elements of the score by the user's relative weights. After each agent navigates five routes, the MoverManager orders the agents from lowest score to highest, then applies the genetic algorithm.

Figure 13 shows a sample scoring comparison. The two agents chose different routes due to different decision vectors. Assume the user set the following values for scoring:

Distance = 0.25

Elevation change = 0.5

Time = 0.75

Security = 1

This means that security is the most important aspect, and distance the least. Additionally, assume the means and standard deviations are as follows:

Mean Distance = 10m, SD = 0m

Mean Elevation change = 15m, SD = 8m

Mean Time = 20 sec, SD = 4 sec

Mean Security Score = 10, SD = 10

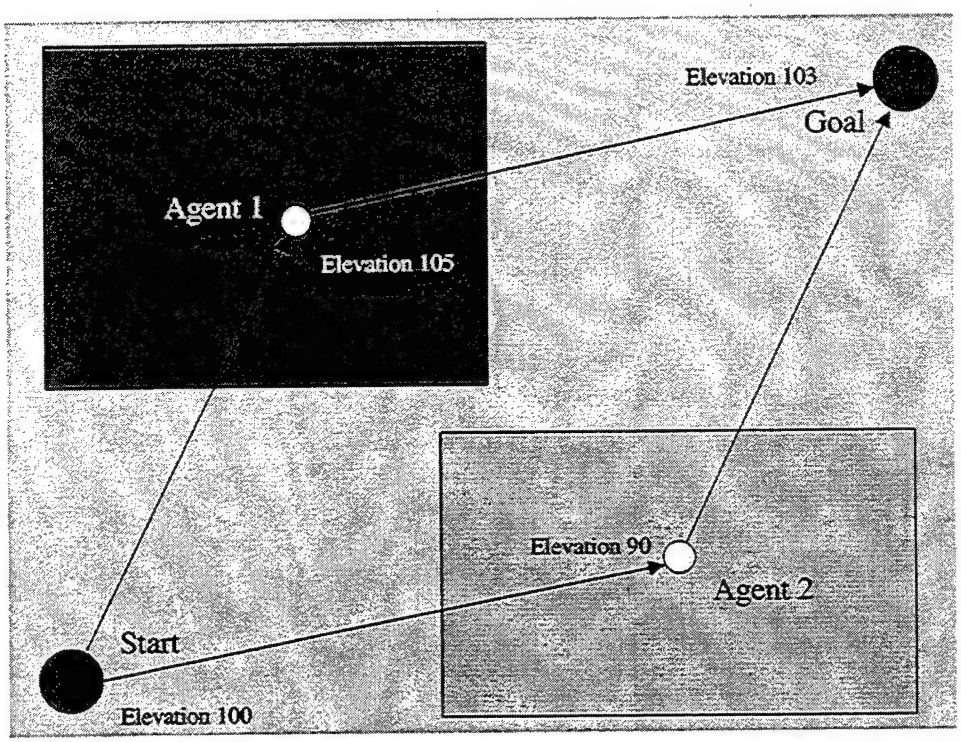

Figure 13. Sample Scoring Calculation

Agent 1 has a total distance traveled of 10m (5m per step, two steps), total elevation change of 7m (5m in the first step, 2m in the second), time elapsed of 24.375 sec (7.5 * 2 for the first step, 7.5 * 1.25 for the second), and a security score of 0. Agent 1's distance is exactly the mean, for a score component of 0. The other score components, expressed as standard deviations from the respective means, are 1.0 for elevation change, -1.1 for time elapsed, and 1.0 for security. Agent 2 has a total distance traveled of 10m, total elevation change of 23m, time elapsed of 16.875 sec, and a security

score of 10 (penalty for crossing an open area). Therefore, Agent 2's score components are 0 for distance, -1.0 for elevation change, 0.78 for time elapsed, and 0 for security. Thus, Agent 1 gets a score of

$$(0)(.25) + (1.0)(.5) + (-1.1)(.75) + (1.0)(1) = 0.68$$

while Agent 2 gets a score of

$$(0)(.25) + (-1.0)(.5) + (0.78)(.75) + (0)(1) = 0.09$$

Therefore, Agent 1 ranks higher than Agent 2 on this route. The MoverManager adds each agent's total scores of five such routes before applying the genetic algorithm. This is a greatly simplified example, as actual routes are much longer and the terrain is more varied. However, the concepts are the same.

2. Genetic Algorithm

The top five agents are the only ones allowed to reproduce to form offspring. These offspring replace the bottom six agents for the next five routes. The top-ranked agent reproduces four times, the second-ranked three times, the third and fourth twice, and the fifth once.

The MoverManager determines the cutoff point for decision vector transfer by random draw. The cutoff point determines how many decision vector elements come from each parent. For example, if the cutoff point is four, four elements come from the first parent, and two come from the second. Additionally, "mutations" occur at a rate of one percent. In these cases, the element at the cutoff point is not taken from either parent, but determined by random draw.

More variation in the decision vectors helps the genetic algorithm change the agent pool. Too much similarity in the top agents causes uniformity, as similar agents

will produce offspring very much like themselves. For this reason, it is very important that all agent decision vectors be generated by random draw as the agent is created. In other words, the user must not attempt to tailor the decision vectors to fit a preconceived notion of success.

3. Achieving Expertise

Since individual agents do not have access to their scores, no learning takes place during execution. All decision vector modification, and thus all expertise development, is done by the MoverManager through the genetic algorithm. The MoverManager continues to apply the genetic algorithm every five routes until 100 routes are complete or the same agent is top ranked three consecutive times, whichever comes last. This prevents an initially successful agent from stopping the genetic algorithm prematurely, and guarantees at least twenty generations of navigators before the system stops adapting. The top-ranked agent then becomes the expert route planner. The system prompts the user for a start and end point, and the expert agent plans the route between them.

THIS PAGE INTENTIONALLY LEFT BLANK

VI. CONCLUSIONS AND RECOMMENDATIONS

A. APPLICATION OF THE RPD MODEL

1. CDM Implementation

CDM proved to be a very effective technique for knowledge elicitation. Although the planned protocol was not followed to the letter, the important information came out in a relatively coherent manner. The timeline proved to be more of a distraction than a discussion aid, and the patrol sketch was the most useful tool. Unfortunately, this makes it difficult to return to the data at a later date for further analysis. The nature of the sketch makes it easy to understand during the interview, but without detailed notes, a rough sketch can be vague or misleading days or weeks later. In the future, videotaped interviews should be the standard.

Data from experts reinforces the initial belief that NDM theory accurately describes expert navigator decision making under the conditions of the experiment. Interviews indicate that intuitive decisions, made quickly and dynamically, are what separate expert navigators from novices. However, the nature of the task makes it difficult to conduct direct observation to corroborate the interview data.

2. Identifying Experts

The single most important, and difficult, aspect of this project was location of true experts as defined by the Dreyfus model. Although the U.S. Army has tens of thousands of infantrymen, a very small percentage of these are truly expert navigators. Many more soldiers consider themselves expert, and some of these are very vocal. True expert navigators are so rare that most infantrymen never serve with one. Consequently, they confuse proficiency, or even competence, with expertise. The capabilities of the seven

experts interviewed for this project are far beyond anything I have personally observed in over nine years of experience in infantry units. A clear definition of expertise, as supplied by our focus group, presented to a group of evaluators who work with an already select group, like the Special Forces instructors, who in turn select only the best performers is one way to get at these rare individuals. Other units and organizations in the military, or more specifically the Special Operations community, may have a comparable subject pool. Conventional line units, however, are unlikely to include even one expert, let alone enough for a study.

B. EXECUTABLE REPRESENTATIONS

1. Suitability of Agent-based Models

Agent-based models provide interesting contrasts to traditional models. Their lack of dependence on rule bases makes modification quick and simple. A decision vector can be modified quickly and easily, as can several system parameters. Unfortunately, this simplicity also means that the system performs slowly.

In test runs, some routes take 30 minutes to execute with 20 agents. The average route takes about six to eight minutes, leading to a total warm-up time of over twelve hours on average on a 350 Mhz processor with 64MB RAM. Once an agent achieves expertise, its average time for execution is under one minute.

Another disadvantage of the agent model is that if the user changes the relative scoring weights, the agent must develop expertise all over again. The refined decision vector is tuned to a specific combination of scoring weights, and will not function properly with a different set.

2. Terrain Representation Techniques

Although the OpenFlight format is widely used and exports directly to many other formats, it does not easily represent data other than elevation and geometry. However, the user has access to the file structure and can take advantage of it. Additional information, like vegetation, roads, and water, is not easy to represent or to access.

The key to this model is grouping polygons by type of terrain they represent. This is a tedious process, even on a small model, and Multi-Gen Creator has no function to automatically organize. The Creator export utility creates VRML files with the same organizational structure, however, and this allows Java classes to determine terrain type by the location of polygons in the file.

3. Portability

VRML allows the model to be viewed on the Internet, and Java classes allow various operating systems to run it. This makes the model easy to access from almost any network, and displayable on any Netscape browser with a VRML plug-in. Systems that do not require the visualization can use the expert agent as well. All calculations are done in the Java classes, so the results could easily be routed into a larger model or system, instead of to a VRML scene graph for display.

C. TRAINING STRATEGIES

1. Training vs. Experience

Most of the senior leaders and instructors that participated in this project expressed skepticism that any training method could speed the transition from novice to expert. They felt that only long and varied experience can produce good navigators, and that attempts to augment or replace current methods are doomed to failure. This

sentiment must be overcome for any new training method or device to be effective. Carefully designed training strategies, proved with scientifically rigorous experiments to be effective, will accomplish this goal. Such strategies or devices need not be computer-based, and should not be touted as substitutes for experience.

It may be possible to train each of the four skills (detailed mental maps, blending cues, dynamic adjustment, and spatial visualization) which contribute to expertise as components of navigation. Further research may help identify other skills, and other feasible training strategies, which may also help speed the transition from novice to expert. The infantry community needs training strategies to develop each of these skills individually, then use navigation courses and training exercises to provide the experience to tie all four together.

Although it seems that experience can be augmented, it will probably never be replaced. No training system, for any reasonably complicated task, is sufficient to produce expertise in a laboratory environment. Navigators need to navigate in the woods, under pressure, in a hostile environment to gain proficiency. Training systems should aim to make the time spent navigating more productive, so that beginners spend more time developing skills, and less time hopelessly lost.

Every training strategy must consider the practical limitations on military training resources. If the Army had plenty of experts who were superb instructors, and provided them with the requisite time for training apprentice navigators, it probably would have no need to consider other alternatives. But the reality of the situation is that there are few experts and, of them, even fewer are also expert instructors. The units who do have these experts certainly are unable to maximize the contact time between them and the students.

Cast under the light of these practical constraints, virtual apprenticeships make more sense.

2. Use of Computer Based Trainers

The executable model presented in Chapter V has two distinct training applications. First, it could be used as an expert advisor to help soldiers learn to plan routes. Exposing more of the agent decision process may help beginners see why one direction is more advantageous than another, and how the competing goals can affect decisions. Second, the expert route planner could be incorporated into higher-level simulations to replicate elite units. Any automated forces that need to move as well-trained, experienced units would move, and that are of sufficiently high resolution to make squad routes useful, could use the model to control movements.

Clearly, some skills lend themselves to computer-based training more than others. It is difficult to imagine a computer-based trainer to help novices develop detailed mental maps more effectively than paper maps and terrain models currently do. Spatial visualization, however, may be a good task to train in a virtual environment. A highly detailed terrain database, with a viewpoint that moves realistically through the environment, coupled with immersive display devices, may help novices to relate map information to real world information. Perhaps a map that looks like a standard paper military map, that morphs to a highly realistic virtual environment of the same area would serve this purpose.

3. Other Possible Trainers

Some of the seven expert subjects worked much harder to achieve expertise than the others. Differences in natural aptitude and quality of training helped some to progress

faster and others more slowly. In this study, the experts who had to work harder to achieve expertise could better verbalize what they do and how they do it. These experts, who understand far better how their specific expertise developed, may provide insight and ideas on how to better train the specific subtasks with conventional tools like orienteering courses or paper map exercises. Such low cost innovations could prove superior to expensive software development projects and high-tech equipment. Further study must be conducted to elicit these types of training devices and strategies.

D. FUTURE WORK

1. Cognitive Task Analysis

Although the cognitive task analysis produced a workable model of expertise, it needs to be validated with direct observation and a larger subject pool. This project abandoned direct observation as too intrusive a method for knowledge elicitation, but as a validation technique it may be more useful. Walking with expert-led patrols to observe and record the expert's actions, not elicit his cognitive processes, should be far less intrusive. The danger is that researchers can waste enormous amounts of time following non-experts. Thus, the interviews will still be a necessary step, if for no other reason than to confirm that the navigator is in fact an expert.

Interviews need to strive for greater depth. More information on how the four key skills develop, exactly how each expert implements them, and whether or not any others exist would improve the model.

2. Improved Agent-based Route Planner

Code efficiency improvements and use of volume rendered terrain in the agent-based route planner would make the route planner far more useful. The twelve hour (or

more) warm-up time could certainly be improved by more efficient Java code, or possibly by translating the code to C++. This latter option, however, would eliminate the visualization portion, since VRML is not compatible with C++. The model contains several large data structures that must be traversed several times in each run. Optimizing one or more of these should provide better performance.

Volume rendered terrain would solve some of the organizational problems caused by the OpenFlight format. Volume rendering uses software constructs known as voxels to represent very small, uniformly sized pieces of terrain. Voxels can store their own terrain types internally, along with other useful information for agent route planning. Such a system requires more memory to store terrain information, but far fewer data structures to store additional information. A single data structure, from which all necessary information can be obtained in a single traversal, obviously would be much more efficient than the current model.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX A. SMALL UNIT MISSIONS

1. GENERAL INFORMATION

Operations studied for this project were basic infantry missions conducted by squads of ten to fifteen men. All closely adhered to standard U.S. Army doctrine since they were conducted as parts of two very rigorous training courses, the Ranger Course and the Special Forces Qualification Course. Two or more instructors accompanied each patrol to evaluate, critique, and ensure the safety of the patrol. Typically, each mission consists of a planning phase, a movement phase, and an actions on the objective phase.

2. PLANNING

Each mission begins with receipt of an order from the instructor, simulating an order from higher headquarters in a real-world mission. Military orders are organized into five sections and are highly standardized (Ranger Handbook). The patrol leader receives a paper copy of the order and is given the opportunity to ask questions. In the Ranger/Special Forces environment, the next four to eight hours are devoted to producing a patrol order, preparing mission-essential equipment, and rehearsing for the operation.

The patrol leader cannot possibly complete all required tasks by himself, so he assigns specific elements of the planning to other members of the patrol. The assistant patrol leader usually writes the supply and maintenance section of the plan, and oversees resupply and maintenance of key systems during the entire planning phase. The radio operator writes the command and signal portion of the plan and ensures all communications equipment is ready for the mission. Other members of the patrol perform equipment maintenance and assist as needed. The patrol leader personally writes

the execution portion of the plan, and his primary assistants are his team leaders. The primary navigator, or "compass man", normally plans the routes and constructs two terrain models of the operational area: one for the movement phase and a smaller, more detailed version for the actions on the objective. The patrol leader spends most of his time on the actions on the objective, and plans backwards from the "hit time", when the (hopefully) first shot is fired, to the time the patrol leaves its planning location. After the plan is written and issued, the patrol rehearses actions on the objective in as much detail as possible. The ideal rehearsal includes all members of the patrol on terrain similar to the objective. The remainder of the available time is spent inspecting mission essential equipment and preparing to move.

3. MOVEMENT

The movement phase begins when the point man leaves the planning area to lead the patrol along its route. The compass man selects the actual route, periodically sending reports back to the patrol leader, who travels near the middle of the formation. The patrol leader changes formations and movement techniques as necessary along the route.

Approximately 400 meters or one major terrain feature away from the objective, the patrol establishes an Objective Rally Point (ORP), where final preparations for actions on the objective take place. First, the patrol halts approximately 150 meters from the planned ORP location. The patrol leader takes a small element forward to confirm the ORP location and its suitability. A good ORP location is covered and concealed, away from roads and trails, and away from suspected enemy locations. If the planned site meets these criteria, the patrol leader moves back to the main body and moves the entire patrol forward to occupy the ORP.

The patrol will ideally spend about one hour in the ORP, but if it has encountered difficulties during the movement phase this time is compressed to the time available. Generally, the patrol leaves all equipment not required for the actual mission in the ORP.

4. ACTIONS ON THE OBJECTIVE

After the ORP is established, the patrol leader takes his team leaders and two security personnel forward for the leaders' reconnaissance of the objective, the first event in the actions on the objective phase. The goal of the leaders' recon is to pinpoint the objective's location and confirm that the plan is workable on the actual terrain. While the leaders are gone, the assistant patrol leader is in charge in the ORP, where the rest of the patrol prepares special equipment, such as demolitions, night vision devices, and obstacle breaching gear, for the mission.

When the leaders return, the patrol leader makes any necessary adjustments to the plan, gives final instructions, and confirms the hit time. The goal is to get all elements in position just in time to execute the operation. The patrol does not want to arrive at the objective too soon, as it risks discovery by the enemy. Obviously, one or more elements not in position at hit time is even less desirable, so the patrol will usually accept sitting in position near the objective for up to ten or fifteen minutes. Generally, the patrol will break into security, support, and assault elements for actions on the objective, and these elements may leave the ORP separately and be separated by several hundred meters during the operation. The actual tasks of each element are totally dependent on the type of mission (e.g. ambush, raid, deliberate attack, defense) and the patrol leader's plan.

At hit time, the patrol leader initiates the operation. He may use a visual signal, radio transmission, predetermined time, or a simple tap on the shoulder, but in any case

the first indication the enemy gets of the patrol's presence in the area should be the firing of the patrol's most powerful weapon. The patrol conducts the mission according to the plan or as instructed by leaders who adjust to changing circumstances. After the mission is complete, the patrol gathers intelligence from the enemy, reconfigures its equipment for movement, and evacuates its casualties. Depending on instructions from superiors, the patrol may remain on the objective or move to a new location. In the Ranger and Special Forces courses, at this point the instructors call the patrol together and conduct a critique of the entire mission.

LIST OF REFERENCES

Dreyfus, H., *Intuitive, Deliberative, and Calculative Models of Expert Performance,* in Zsambok, C., and Klein, G. (Eds.), *Naturalistic Decision Making*, Lawrence Erlbaum, 1997.

Ferber, J., *Multi-agent Systems,* Addison Wesley Longman, 1999.

Hoffman, R. R., and others, "Eliciting Knowledge from Experts: A Methodological Analysis," *Organizational Behavior and Human Decision Processes,* 62 (1), 1995.

Holland, J. H., *Hidden Order: How Adaptation Builds Complexity,* Perseus Books, 1995.

Ilachinski, A., *Irreducible Semi-Autonomous Adaptive Combat (ISAAC): an Artificial Life Approach to Land Combat,* CRM97-61.10, Center for Naval Analyses, 1997.

Klein, G., *Sources of power: How people make decisions*, MIT Press, 1998.

Klein, G., Calderwood, R., and MacGregor, D., "Critical Decision Method for Eliciting Knowledge," *IEEE Transactions on Systems, Man, and Cybernetics*, 19 (3), 1989.

Klein, G., and others, "Characteristics of Skilled Option Generation in Chess," *Organizational Behavior and Human Decision Processes,* 62 (1), 1995.

National Imagery and Mapping Association, *Guide to NIMA Products and Services,* Government Printing Office, 1999.

Orasanu, J., and Connolly, T., *The Reinvention of Decision Making,* in Klein, G., and others (Eds.), *Decision Making in Action: Models and Methods,* Ablex, 1993.

Peterson, B., Stine, J., and Darken, R., "Eliciting Knowledge from Military Ground Navigators," paper presented at 5th Conference on Naturalistic Decision Making, Stockholm, Sweden, 2000.

U.S. Army Infantry School, *Ranger Handbook,* Department of the Army, 1991.

Zsambok, C., *Naturalistic Decision Making: Where Are We Now?,* in Zsambok, C., and Klein, G., (Eds.), *Naturalistic Decision Making*, Lawrence Erlbaum, 1997.

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center 2
   8725 John J. Kingman Road, STE 0944
   Ft. Belvoir, Virginia 22060-6218

2. Dudley Knox Library 2
   Naval Postgraduate School
   411 Dyer Rd.
   Monterey, California 93943-5101

3. Commander 1
   4th Ranger Training Battalion
   Ft. Benning, Georgia 31905

4. Commander 1
   Company F, Special Warfare Instruction Center
   Ft. Bragg, North Carolina 28307

5. Dr. Rudolph Darken 1
   MOVES Academic Group
   Naval Postgraduate School
   Monterey, California 93943-5221

6. Mr. Barry Peterson 1
   MOVES Academic Group
   Naval Postgraduate School
   Monterey, California 93943-5221

7. Mr. John Hiles 1
   MOVES Academic Group
   Naval Postgraduate School
   Monterey, California 93943-5118

8. LCDR Dylan Schmorrow 1
   Office of Naval Research
   800 N. Quincy St. Tower 1
   Arlington, Virginia 22217-5660

9. MAJ Jason L. Stine 2
   603 East 35th Avenue
   Winfield, Kansas 67156